Advanced Simulation and Computing

FY15–19 Program Notebook

for
Advanced Technology Development & Mitigation, Computational Systems and Software Environment
and
Facility Operations and User Support

NA-ASC-128R-14-VOL.2-PN
LLNL-TR-665217

September 22, 2014

Prepared for the tri-lab ASC Program by:

Lawrence Livermore National Laboratory is operated by Lawrence Livermore National Security, LLC, for the U.S. Department of Energy, National Nuclear Security Administration under Contract DE-AC52-07NA27344.

Disclaimer

Contents

Advanced Technology Development and Mitigation (WBS 1.5.6)

The Advanced Technology Development and Mitigation (ATDM) sub-program includes laboratory code and computer engineering and science projects that pursue long-term simulation and computing goals relevant to the broad national security missions of the National Nuclear Security Administration (NNSA).

Next-Generation Architecture and Software Development (WBS 1.5.6.2)

This product is focused on long-term computing technology research to influence the shift in computing technology to extreme, heterogeneous architectures and to mitigate its impact and advance its capabilities for Advanced Simulation and Computing (ASC) simulation codes.

Next-Generation Computing Enablement and Co-Design (LLNL)

The Next Generation Computing Enablement and Co-Design efforts will prepare ASC for the next generation of advanced computing technologies beyond the current advanced technology system (ATS) projects. This project addresses both the software environment and future architectures. It includes coordinating next-generation activities with Integrated Codes (IC) at Lawrence Livermore National Laboratory (LLNL), within the tri-lab, and externally. The efforts will support IC ATDM next-generation application efforts to prepare for new platforms and to adapt other codes to the expected new architectures. The software efforts include system level software, resource management, development tools, data analysis tools, parallel input/output (I/O), and programming models. On the hardware side, these efforts include tracking and collaborating on technology innovations. This effort includes interactions with the Department of Energy's (DOE's) Advanced Scientific Computing Research (ASCR), vendors, and academia, including planning and technical coordination for vendor contracts. Team members will carry out investigations and co-design activities using test beds and existing technology, making use of proxy applications.

Required Capabilities

R1: Understanding challenges the ASC Program must overcome to field next-generation systems and next-generation applications

R2: Prototypes and simulators to assess the impact on ASC applications

Five-Year Plan

FY15

- Conduct co-design activities with co-design centers and vendors, and evaluate next-generation technologies
- Research and develop tools and techniques to effectively utilize emerging next-generation architectures, including tools for measuring impact on power, energy, and performance of applications; next-generation resource management framework; performance modeling; and techniques for input/output (I/O) stack optimizations for future parallel file systems
- Evolve the development environment to include next-generation debugging and to support programming models used in IC
- Participate in the technical management of DesignForward projects
- Execute LLNL portion of the tri-lab Level 2 milestone to demonstrate advances in proxy applications through programming abstractions or performance gains
- Procure and deploy prototypes and leverage test beds to allow exploration of how emerging technologies support ASC applications

FY16

- Conduct co-design activities with vendors, and research and evaluate next-generation technologies
- Research and develop tools and techniques to effectively utilize emerging next-generation architectures (for example, next-generation resource management; programming models; tools for measuring impact on power, energy, and performance of applications; future I/O stack optimizations; next-generation debugging; and performance modeling)
- Coordinate and administer FastForward2 contracts, and participate in DesignForward
- Procure and deploy prototypes and test beds to allow exploration of how emerging technologies support ASC applications

FY17

- Conduct co-design activities with vendors, and research and evaluate next-generation technologies
- Research and develop tools and techniques to effectively utilize emerging next-generation architectures (for example, next-generation resource management; programming models; tools for measuring impact on power, energy, and performance of applications; future I/O stack optimizations; next-generation debugging; and performance modeling)
- Exercise applications on test beds to identify gaps and enhancements needed for next-generation IC codes
- Coordinate and administer PathForward contracts, such as FastForward and DesignForward
- Identify technology gaps in available systems and procure and deploy prototypes and test beds to allow exploration of how those technologies support ASC applications

FY18

- Exercise applications on test beds to identify gaps and enhancements needed for next-generation IC codes
- Identify gaps and enhancements needed for technologies in co-design effort with ASC IC
- Coordinate and administer contracts that lay the groundwork for the next-generation machines beyond the current ATS
- Identify technology gaps in available systems and procure and deploy prototypes and test beds to allow exploration of how those technologies support ASC applications

FY19

- Explore advanced technology developments

- Administer contracts for next-generation research in software environments
- Conduct co-design activities with IC, other labs, and vendors
- Exercise applications on test beds or current advanced architectures to identify gaps and enhancements needed for next-generation IC codes

FastForward—Industrial Partnerships for Extreme-Scale Technology Research and Development (LLNL)

The FastForward program is a jointly funded collaboration between DOE Office of Science and NNSA to initiate partnerships with multiple companies to accelerate the R&D of critical technologies needed for extreme scale computing, on the path toward exascale computing. This program is administered by DOE and contracted through Lawrence Livermore National Security, LLC, as part of a seven-lab consortium (Argonne, Lawrence Berkeley, Lawrence Livermore, Los Alamos, Oak Ridge, Pacific Northwest, and Sandia national laboratories).

The first set of five FastForward projects, awarded in summer 2012, included:

- AMD, with focus on processor- and memory-related technologies
- IBM, with focus on memory-related technologies
- Intel, with focus on processor- and memory-related technologies
- Nvidia, with focus on processor and memory-related technologies
- WhamCloud (now owned by Intel), with focus on storage and I/O technologies

Required Capabilities

R1: Commercially available within 5–10 years processor, memory, and storage, and I/O technologies that will maximize energy and concurrency efficiency while increasing the performance, productivity, and reliability of key DOE extreme-scale applications

Five-Year Plan

FY15

- Provide technical coordination and contractual management for FastForward2 contracts; project ends in FY15

Co-Design and Programming Model Research (LANL)

This project contains the forward-looking research for advance computing technologies at extreme-scale. Co-design research and programming model research are the base of these investigations in support of ASC code needs on future hardware.

The co-design component of the project leverages other activities at Los Alamos National Laboratory (LANL) to build a co-design process through the collaborative creation of patterns, strategies, and abstractions for the implementation and optimization of scientific applications and algorithms on emerging hardware architectures. One aspect of this process will be a suite of open-source proxy applications, derived from and feeding back into ASC IC teams. ASC code teams have informed and continue to provide the requirements for the study based upon application domains of interest to ASC. Specification documents and reference implementations produced will act as the basis for most of the work in this project, ensuring that it is targeted directly to ASC code needs.

The project will track the development of next-generation hardware architectures and study both computational and data movement patterns represented by the chosen mini applications (mini-apps) developed in conjunction with ASC IC code developers. A major goal of this effort is to inform application developers of methods and best practices that will be necessary for code development on Trinity-like architectures. LANL will also investigate novel hardware data transformation techniques prototyped using reconfigurable hardware.

Work being done directly with IC teams will support the evolution of current codes towards next-generation architectures by providing computer science expertise on improved mesh data structures, new strategies for adaptive mesh refinement (AMR), performance improvement, data locality, and compressed data structures for materials. In addition, LANL will explore increased software abstraction, through emerging programming models and DSLs.

The programming models aspect of the project studies emerging hardware and software trends and their impact on programming abstractions/models, the overall software development tool chain, and run-time systems for scientific-computing environments. LANL's goal is to develop a set of technologies that will assist in the development of the next-generation of application codes as well as extend the lifetime of current codes.

The impact of new architectures on critical IC codes has been categorized at a high level in terms of removing bulk-synchronous communications and increasing levels of concurrency and parallelism. A critical element in meeting these challenges is the adoption of new approaches to programming that reduce the introduction of these characteristics and simplify the programmability of future systems. In FY15, this project will be responsible for a tri-lab Level 2 milestone with IC that will investigate a potential path forward for reaching these goals by developing and applying these techniques. In addition to improving concurrency and reducing synchronization points, LANL will explore approaches that enable interoperability with message passing interface (MPI)-

based codes to not only minimize the overall impact but also, more importantly, provide a staged migration path for existing codes.

This project also supports an effort to develop and validate discrete-event simulation techniques to the modeling of advanced file systems such as burst-buffer and hierarchical and caching storage systems.

Required Capabilities

R1: Provide representative set of proxy applications that explore focused functional capabilities of the ASC code base (for example, unstructured meshes and particle operations)

R2: Provide co-design knowledge base consisting of recommendations and best practices for code development on Trinity and future extreme-scale architectures

R3: Working with vendors, explore areas where custom hardware can be applied to applications of interest and develop prototypes for demonstration

R4: Design and demonstration of low-level runtime components, for example, lightweight-thread scheduling and memory affinity on both shared and distributed memory architectures

R5: Language and compiler design and development (including memory and process/thread management), focused on needs of IC workloads

Five-Year Plan

FY15

- Extend and enhance next-generation programming techniques in the context of an ASC integrated code
- Extend and support program compiler and analysis tool-chain on next-generation architectures for characterization and analysis of ASC applications
- Characterize and validate proxy applications in support of the tri-lab CSSE/IC Level 2 milestone, DesignForward, and FastForward2 efforts, and in support of preparing Eulerian and Lagrangian application project codes for Trinity
- Extend parallel file-system simulation capabilities for hierarchical file-systems (Burst Buffer) and make it a usable tool for design and validation

FY16

- Work with ASC developers to fully transition dynamic, multiphysics applications to current heterogeneous hardware
- Continue to explore and document the advantages/disadvantages of emerging HPC programming models for IC workloads on emerging architectures; focus on scalability and resiliency implications and support open-source releases
- Continue to design and implement tool chain capabilities

- Work with ASC developers to extend dynamic application to full system, including in situ analysis and fault management using services expected to be provided by vendor-supplied system runtime
- Continue to explore and document the advantages and disadvantages of emerging high performance computing (HPC) programming models for the needs of the IC workloads

FY17

- Work with ASC developers to extend dynamic application to full system, including in situ analysis and fault management using services expected to be provided by vendor-supplied system runtime
- Continue to explore and document the advantages and disadvantages of emerging HPC programming models for the needs of IC workloads
- Continue to design and implement tool chain capabilities

FY18

- Work with ASC developers to exploit dynamic applications on extreme-scale computing platforms
- Continue to explore and document the advantages and disadvantages of emerging HPC programming models for the needs of IC workloads
- Continue to design and implement tool chain capabilities

FY19

- Continue to explore and document emerging and dynamic HPC and non-HPC programming models for the needs of IC workloads
- Continue to design and implement tool chain capabilities for advanced architectures

Programing Models and Abstractions (SNL)

The next generation of computing platforms promises both new capabilities and increased capacity for meeting SNL's mission challenges. However, these platforms will involve new computer architectures and using this new capability effectively will require an investment in new codes and algorithms that can effectively express and exploit multiple levels of parallelism. This project explores programing models and abstractions that will allow our new ATDM codes to effectively utilize next-generation hardware accounting for hybrid parallelism and the need for enhanced resilience. As part of our risk mitigation ATDM strategy, we will also invest and mature transitional models such as Kokkos that not only provide a means of implementing new codes, but can also be used to refactor low-level kernels in our existing IDCs.

Note: this reflects early planning for ATDM and we expect that required capabilities and detailed deliverables will be refined and extended as ATDM planning continues to mature.

Required Capabilities

R1: The ability to insulate the implementation of physics from the details of data-layout and performance implementation

R2: The ability to effectively program for heterogeneous architectures with deep memory hierarchies ranging from non-volatile memory, DRAM, and traditional caches

R3: Programming models that support both data-parallel and task-based parallelism in flexible and interactive ways

Five-Year Plan

FY15:

- Execute SNL's portion of the tri-lab Level 2 milestone *Demonstrate Advances in Proxy Applications Through Performance Gains and/or Performance Portable Abstractions*
- Execute Level 2 milestone "*Programming Models Analysis for Next-Generation Platforms*" by prototyping asynchronous task-based programming models on Mantevo mini-apps
- Summarize and provide lessons learned from prototype mini-apps to ATDM code teams
- Evaluate the performance and usability of Kokkos in proxy applications representative of implicit finite element production codes
- Develop an implementation of an asynchronous many task (AMT) runtime on an advanced technology system
- Develop optimized dynamic scheduling and work-stealing methods for AMT runtime

- Integrate in-memory data storage capabilities (Kelpie) into AMT runtime
- Investigate approaches to reduce data-movement for integrated application workflows

FY16:

- Collaborate with ATDM library and ATDM application teams to prototype strategies to migrate their capabilities to heterogeneous computing architectures
- Research hierarchical, heterogeneous domain decomposition to address NUMA performance concerns and work decomposition in a MPI + CPU + accelerator strategy
- Develop an implementation of a resilient, scalable asynchronous many task (AMT) runtime

FY17:

- Update and summarize lessons learned from prototype mini-apps to SNL code teams
- Develop proxy-apps with deep heterogeneous parallel work hierarchies
- Demonstrate one or more ASC Mantevo applications running at scale on a capability class machines using the AMT runtime
- Develop optimized dynamic scheduling and work-stealing methods for AMT runtime
- Work with ATDM code teams to support AMT runtime within target ATDM codes

FY18:

- Develop proxy-apps to evaluate programming models for processor-in-memory architectures
- Demonstrate AMT runtime on ATDM codes on both test beds
- Further optimized AMT scheduling and work-stealing methods with lessons learned from ATDM codes

FY19:

- Provide technical requirements for new hardware agnostic software stack defining and development and execution environment for portable performance of our codes
- Define requirements for a synergistic task-based parallel runtime and data parallel execution model with optimizes the performance of our code base

DesignForward—Industrial Partnerships for Extreme-Scale Technology Research and Development (LBNL)

The DesignForward program is a jointly funded collaboration between DOE Office of Science (SC) and NNSA to initiate partnerships with multiple companies to accelerate the R&D of critical technologies needed for extreme scale computing, on the path toward exascale computing. This program is administered by DOE and contracted through Lawrence Berkeley National Lab, as part of a seven-lab consortium (ANL, Lawrence Berkeley, LLNL, LANL, ORNL, Pacific Northwest, and SNL). DesignForward seeks to fund innovative new and/or accelerated R&D of technologies targeted for productization in the 5–10 year timeframe.

The five DesignForward Interconnect projects were awarded in November 2013 to the following companies: AMD, Cray, IBM, Intel, and NVIDIA.

Required Capabilities

R1: Commercially available within 5–10 years the needed interconnect technologies that will maximize energy and concurrency efficiency while increasing the performance, productivity, and reliability of key DOE extreme-scale applications

Five-Year Plan

FY15

- Provide technical coordination and contractual management for DesignForward contracts in the Interconnect and System Integration contracts (project ends in FY15)

Computational Systems and Software Environment (WBS 1.5.4)

The mission of this national sub-program is to build integrated, balanced, and scalable computational capabilities to meet the predictive simulation requirements of the NNSA. This sub-program strives to provide users of ASC computing resources a stable and seamless computing environment for all ASC-deployed platforms. Along with these powerful systems that ASC will maintain and continue to field, the supporting software infrastructure that Computational Systems and Software Environment (CSSE) is responsible for deploying on these platforms includes many critical components, from system software and tools, to I/O, storage and networking, to post-processing visualization and data analysis tools. Achieving this deployment objective requires sustained investment in applied research and development (R&D) activities to create technologies that address ASC's unique mission-driven needs for scalability, parallelism, performance, and reliability.

Commodity Technology Systems (WBS 1.5.4.8)

This level 4 product provides production platforms and integrated planning for the overall system architecture commensurate with projected user workloads. The scope of this product includes strategic planning, research, development, procurement, hardware maintenance, testing, integration and deployment, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include strategic planning, performance modeling, benchmarking, procurement and integration coordination, and installation. This product also provides market research for future Commodity Technology Systems (CTS).

Production Planning and Integration (LLNL)

The LLNL ASC strategy for commodity technology systems is to leverage industry advances and open source software standards to build, field, and integrate Linux clusters of various sizes into classified and unclassified production service. The programmatic objective is to dramatically reduce overall total cost of ownership of these commodity systems relative to best practices in Linux cluster deployments today. This objective strives to quickly make these systems robust, useful production clusters under the coming load of ASC scientific simulation capacity workloads.

Required Capabilities

R1: Robust commodity Linux compute clusters to run ASC simulations

Five-Year Plan

FY15

- Successfully award the contract for CTS-1, based on tri-lab CTS-1 process and review
- Plan for deployment of LLNL CTS-1 systems

FY16

- Monitor computer industry developments for opportunities to enhance capacity computing and the associated infrastructure
- Deploy CTS-1 systems

FY17

- Monitor computer industry developments for opportunities to enhance capacity computing and the associated infrastructure

FY18

- Monitor computer industry developments for opportunities to enhance capacity computing and the associated infrastructure
- Plan for procurement of CTS-2 systems

FY19

- Procure CTS-2 systems
- Plan for deployment for CTS-2 systems

Computing Platform Integration and Deployment (LANL)

The scope of the Computing Platform Integration and Deployment project is to accept delivery and begin deployment of production systems. Primary capabilities include completing the acceptance tests, diagnostics tests, integrating the systems into the LANL yellow network, system stabilization, and transition into the classified network.

Required Capabilities

R1: Understanding of industry roadmaps and technologies available in the CTS-1 timeframe

R2: Working with vendors and users, identify potential architectures to meet mission need

Five-Year Plan

FY15

- Continue to operate Luna and the other capacity systems in both the classified and unclassified computing environments
- Provide production support for CTS systems
- Participate in the tri-lab selection for the NNSA ASC CTS-1 procurement

FY16

- Integrate CTS-1 systems into the LANL HPC computing environments
- Provide production support for CTS-1 systems
- Decommission TLCC2 systems

FY17

- Provide production support for CTS-1 systems

FY18

- Provide production support for CTS-1 system
- Participate in the tri-lab selection for the NNSA ASC CTS-2 procurement

FY19

- Plan for the deployment of CTS-2 systems

ASC Commodity Systems (SNL)

The purpose of the ASC Commodity Systems project is to support the acquisition, delivery, and installation of new ASC commodity technology systems. The project is supported by analysis of SNL's portfolio of application needs for capacity workload systems within the context of the higher integrated ASC platform strategy of commodity and advanced technology systems. Efforts include definition of requirements for TLCC systems and collaboration with the Common Computing Environment (CCE) product, with respect to a common software stack for new and existing capacity systems.

Required Capabilities

R1: Provide production computing cycles for ASC applications, especially those that for technical or programmatic reasons are not expected to run on advanced technology systems

Five-Year Plan

FY15

- Participate in CTS-1 selection and contract reviews

FY16

- Accept delivery and install new CTS-1 system(s)

FY17

- Plan for and execute transition of TLCC-2 systems to institutional use and drop from ASC budget

FY18

- Monitor industry hardware and software roadmaps to understand opportunities for cost-effective integration of new commodity technology into future CTS and implications or requirements for future CCE software development

FY19

- Evaluate need for technology refresh of CTS-1 systems (processor or memory)
- Participate in CTS-2 planning activities

Advanced Technology Systems (WBS 1.5.4.3)

This level 4 product provides advanced architectures in response to programmatic, computing needs. The scope of this product includes strategic planning, research, development, procurement, testing, integration and deployment, as well as industrial and academic collaborations. Projects and technologies include strategic planning, performance modeling, benchmarking, and procurement and integration coordination. This product also provides market research, and the investigation of advanced architectural concepts and hardware (including node interconnects and machine area networks) via prototype development, deployment, and test bed activities. Also included in this product are cost-effective computers designed to achieve extreme speeds in addressing specific, stockpile-relevant issues through development of enhanced performance codes especially suited to run on the systems.

Sequoia Tri-Lab Advanced Technology Platform (LLNL)

Sequoia is a 20-petaFLOP/s IBM BlueGene/Q (BG/Q) ATS platform that was sited at LLNL in FY12 with final acceptance in early FY13. BG/Q brings many innovations over the previous BG generations, including 16 cores per node, multithreaded cores, a five-dimensional torus interconnect, water cooling, and optical fiber links. The 20-petaFLOP/s system has a staggering 1.6 million processor cores with a total possible 102 million hardware threads all operating simultaneously. This type of parallelism dictates new directions in supercomputing and enters a new regime of the possible physical systems that can be simulated numerically. Codes that are optimized for multi-core and multi-threading will run best on this machine. This platform will be used as a Capability Computing Campaign (CCC) machine for tri-lab stockpile stewardship milestones. Every six months a new CCC process will be run and the next suite of codes will be ushered onto the machine.

Required Capabilities

- R1: Production weapons science calculations at scale (72K nodes and less)
- R2: Production codes as part of CCC process

Five-Year Plan

FY15

- Run two CCC processes
- Continue to investigate optimal performance tuning for specific codes

FY16

- Run two CCC processes
- Continue to investigate optimal performance tuning for specific codes

FY17

- Run two CCC processes
- Continue to investigation optimal performance tuning for specific codes

FY18

- Retire Sequoia after bringing in 2017 system; project ends

Sierra Tri-Lab Advanced Technology System (LLNL)

To support the ongoing and coming workload for the stockpile with the necessarily powerful computing systems, the NNSA ASC Program is requesting authorization and funding for a classified uncertainty quantification (UQ)- and weapons-science-focused ATS, to be sited and integrated in 2017, that will fill a critical role in support of the Directed Stockpile Work (DSW) mission during the FY18–FY22 period. The Sierra ATS will replace Sequoia and its mission, which will be past its useful lifetime by FY18.

The platform procurement is called "the Sierra ATS procurement," and the implementation project is called the "Sierra ATS integration project." While it will be sited at LLNL, the operation will fall under a proven national user facility paradigm and the system will be available to LANL, LLNL, and SNL.

In keeping with the mission requirement to field an ATS, some portion of the Sierra procurement budget (~10–15%) will be devoted to non-recurring engineering (NRE) work in partnership with the selected vendor.

LLNL is partnering with two Office of Science Labs, Argonne National Laboratory (ANL) and Oak Ridge National Laboratory (ORNL), to acquire three leadership computing systems, one being the Sierra ATS. This collaboration is called CORAL (Collaboration of Oak Ridge, Argonne, and Livermore).

Required Capabilities

- R1: Capable of running an uncertainty quantification throughput workload on a suite of integrated weapons performance calculations of relevance and importance to the Stockpile Stewardship Program (SSP) (primary mission)
- R2: Capable of running a scalable science workload on selected calculations relevant to supporting the predictive capability framework goals (primary mission)
- R3: Capable of running very large benchmark weapon simulations and non-nuclear assessments (secondary mission)
- R4: Accelerate platform innovation by pursuing new technology paths

Five-Year Plan

FY15

- Provide technical coordination and contractual management for CORAL non-recurring engineering and Sierra contracts
- Begin application preparations for Sierra system through Center of Excellence

FY16

- Provide technical coordination and contractual management for CORAL NRE and the Sierra ATS contracts

- In collaboration with ANL and ORNL, work with the two NRE contract winners to reduce identified gaps
- In collaboration with ANL and ORNL, evaluate the NRE results, make the Go/No-Go decision, and negotiate the replacement of target requirements with hard requirements in the CORAL build contracts
- Complete FY16 facility power modernization

FY17

- Provide technical coordination and contractual management for the Sierra ATS contracts
- Complete computer room upgrades for siting the Sierra ATS
- Accept delivery of the Sierra ATS

FY18

- Transition the Sierra ATS to the classified network

FY19

- Run two CCC processes for Sierra
- Continue to investigate optimal performance tuning for specific codes

Hyperion Test Bed (LLNL)

With the extreme demands for capacity computing, the I/O requirements of petascale applications for Sequoia, and the need for improved scientific data management capabilities, it is clearly apparent that emerging breakthrough technologies need to be tested in a large-scale environment such as Hyperion. The Hyperion Test Bed project will work with an expanded set of Hyperion vendor partners in the next phase of the Hyperion project to evaluate innovative node architectures, networks, and alternative storage solutions. Hyperion will continue to be a unique and critical resource for the functionality, performance, stability, and scalability testing of system software.

Required Capabilities

R1: Open, large-scale testing capability for current and future commodity technology systems I/O, computing, and networking technologies

Five-Year Plan

FY15

- Procure and deploy a technology refresh of the Hyperion phase 2 servers, including high performance cluster interconnect, and potentially storage class memory
- Explore new system software models to evaluate the use of high performance storage class memory and the design impacts of storage class memory on future system software and hardware architectures
- Continue to support scalability testing on system software, middleware, storage, and file systems

FY16

- Continue to support scalability testing on system software, middleware, storage, and file systems
- Complete deployment of Hyperion phase 1 server technology refresh, including high performance storage class memory and data intensive capabilities
- Procure and begin deployment of a technology refresh for the Hyperion phase 2 servers, including high performance storage class memory, cluster interconnect, and data intensive capabilities
- Explore new software models to evaluate the use of high performance storage class memory and the design impacts of storage class memory on future system software and hardware architectures

FY17

- Continue to support scalability testing on system software, middleware, storage, and file systems

- Complete deployment of Hyperion phase 1 server technology refresh, including high performance storage class memory, and data intensive capabilities
- Procure and begin deployment of a technology refresh for Hyperion phase 2 servers, including high performance storage class memory, cluster interconnect, and data intensive capabilities
- Explore new software models to evaluate the use of high performance storage class memory and the design impacts of storage-class memory on future system software and hardware architectures

FY18

- Continue to support scalability testing on system software, middleware, storage, and file systems
- Complete deployment of Hyperion phase 2 server technology refresh, including high performance storage class memory, cluster interconnect, and data intensive capabilities
- Procure technology refresh for Hyperion parallel file system storage and storage area network
- Explore new software models to evaluate the use of high performance storage class memory and the design impacts of storage class memory on future system software and hardware architectures

FY19

- Continue to support scalability testing on system software, middleware, storage, and file systems
- Explore and evolve system software models for high-performance storage-class memory and the design impacts of storage-class memory on future system software and hardware architectures

Advanced Architecture Test Bed Research and Development (SNL)

This project will address a critical need for a range of experimental architecture test beds to support path-finding explorations of alternative programming models, architecture-aware algorithms, low-energy runtime and system software, and advanced memory subsystem development. The systems will be used to develop Mantevo proxy applications, enable application performance analysis with Mantevo proxy applications, support the Heterogeneous Computing and Programming Model R&D, the Software and Tools for Scalability and Performance projects, and for SST validation efforts. These test bed systems are made available for "test pilot" users who understand the experimental nature of these test beds. At the present time, it is more important to explore a diverse set of architectural alternatives than to push large scale. Discussions continue with Intel, Advanced Micro Devices (AMD), IBM, NVIDIA, Micron Technology, and other computer companies regarding ASC interest in obtaining early access to experimental architecture test beds. These partnerships will establish a strong foundation for co-design activities that can influence future hardware designs.

Required Capabilities

R1: Provide a range of advanced system technology experimental architecture platforms for the investigation and advancement of programming model and algorithmic experimentation of ASC applications

R2: Collaborate with industry on lessons learned and opportunities for improvement as pre-production technology matures into production capabilities

R3: Support transition of stable advanced architecture test beds into IC computing environments to support testing and evaluation with ASC production codes

R4: Support ASC Predictive Science Academic Alliance Program II (PSAAP II) and ASCR lab access to SNL's Advanced Architecture test beds

Five-Year Plan

FY15

- Provide platforms to support investigating new programming models and evaluating compilers and application performance (including existing system technology refresh and network relocation based on program needs)
- Provide platforms and/or devices for advanced power and energy research and in situ application power and energy analysis
- Provide platforms for advanced node and platform-level architecture analysis and investigations supporting next-generation platforms

FY16

- Update platforms to support investigating new programming models, and evaluating compilers and application performance (including existing system technology refresh and network relocation based on program needs)

- Update platforms and/or devices for advanced power and energy research and in situ application power and energy analysis
- Update platforms for advanced node and platform-level architecture analysis and investigations supporting next-generation platforms

FY17

- Update platforms to support investigating new programming models, and evaluating compilers and application performance/resilience (including existing system technology refresh and network relocation based on program needs)
- Update platforms and/or devices for advanced power and energy research and in situ application power and energy analysis to support dynamic power management
- Update platforms for advanced node and platform-level architecture analysis and investigations supporting next-generation platforms; expand to support prototype memory and interconnect technologies

FY18

- Update platforms to support investigating new programming models, and evaluating compilers and application performance/resilience (including existing system technology refresh and network relocation based on program needs)
- Update platforms and/or devices for advanced power and energy research and in situ application power and energy analysis to support dynamic power management
- Update platforms for advanced node and platform-level architecture analysis and investigations supporting next-generation platforms; expand to support prototype memory and interconnect technologies

FY19

- Update platforms to support investigating new programming models, and evaluating compilers and application performance/resilience (including existing system technology refresh and network relocation based on program needs)
- Update platforms and/or devices for advanced power and energy research and in situ application power and energy analysis to support dynamic power management
- Update platforms for advanced node and platform-level architecture analysis and investigations supporting next-generation platforms; expand to support prototype memory and interconnect technologies

Application Performance Analysis (SNL)

The purpose of the Application Performance Analysis project is to develop tools, techniques and methodologies to support the analysis and evaluation of current and next-generation HPC technologies. A primary focus area of the project is to provide leadership for the Mantevo[1] project, and to facilitate the use of Mantevo mini and proxy applications ~~use~~ as a tool. In addition, the project will utilize classic empirical and mathematical performance analysis methods to achieve its goals.

Next-generation computing platforms are expected to present significantly different architectural designs from previous platforms. In preparation for these changes, the project will explore the potential computing environments from processor ~~core~~, to node, to inter-node. Mini and proxy applications, test beds, simulation capabilities provided by the SST, abstract machine models, and analytic performance models will be used. The outcome will be a better understanding of the characteristics and capabilities within the context of the computational science and engineering simulations of interest to the ASC program on emerging and future architectures and will inform hardware and software requirements. A primary activity for FY14 is to study and identify key performance issues of applications executing on emerging technologies and, in particular, the Trinity ATS.

Required Capabilities

R1: Expert knowledge of ASC applications, algorithms, high performance computer architectures, hardware, and software; ability to translate that expert knowledge to definition of mini and proxy applications

R2: Ability to apply mini/proxy applications and classical performance analysis and modeling techniques to evaluate next-generation technologies and platforms

Five-Year Plan

FY15

- Work with co-design centers and activities that apply to mini and proxy applications
- Apply mini applications in support of the Trinity procurement
- Support Level 2 milestone *Using Performance Modeling and Simulation Tools and Techniques to Gauge Key Application Performance Characteristics of the Trinity Platform*

FY16

- Work with co-design centers and activities that apply to mini and proxy applications
- Apply mini-apps and participate in performance analysis in support of the Trinity procurement

[1] https://software.sandia.gov/mantevo/

- Provide support to the ATS-3 design team in defining the role of mini-apps in the procurement evaluation and acceptance phases of the acquisition

FY17

- Work with co-design centers and activities that apply to mini and proxy applications
- Provide support to the ATS-3 design team in defining the role of mini-apps in the procurement evaluation and acceptance phases of the acquisition

FY18

- Work with co-design centers and activities that apply to mini and proxy-apps
- Work with the selected ATS-3 vendor and the tri-labs in the implementation of mini-apps and capability applications for procurement acceptance

FY19

- Work with co-design centers and activities for proxy applications and architectures
- Provide support in evaluation of the ATS-2 platform

Alliance for Computing at Extreme Scale Trinity Advanced Technology System (LANL, SNL)

The objective of this project is to define requirements and potential system architectures for platforms that meet future ASC programmatic requirements and drivers. The primary activity is to lead the design, acquisition, and plan for deployment of the Trinity ATS. The project will coalesce mission requirements, application algorithms, user requirements, and HPC computer industry hardware/software trends into the design and operation process.

The Trinity platform is the first ATS for the updated ASC computing strategy. The project is a joint collaboration of the New Mexico Alliance for Computing at Extreme Scale (ACES), a partnership between Los Alamos National Laboratory and Sandia National Laboratories. The ACES partners have signed a Memorandum of Understanding (MOU) to perform the following for Trinity and subsequent ATS platforms:

- Provide a user facility for HPC to the NNSA weapons program in support of stockpile stewardship
- Develop requirements and definitions for architectural and system design, and of ASC systems
- Provide procurement support
- Develop key technology
- Provide project management
- Provide systems deployment
- Provide operations
- Provide user support

The architecture and design of Trinity is to provide performance for large-scale applications in support of the NNSA program's most challenging problems. This project covers all aspects of the technical, programmatic, and procurement planning for the platform.

Trinity will replace the Cielo platform sited at LANL but will be used by the applications users at each of the NNSA labs.

Required Capabilities

R1: Architecture, design, and operational support under the ACES MOU

R2: Understanding of tri-lab ASC mission needs and drivers

R3: Understanding of industry roadmaps and the technologies available in the Trinity timeframe

R4: Ability to translate mission needs and vendor capabilities into a platform architecture and associated technical requirements for a successful procurement and deployment process

Five-Year Plan

FY15

- Provide technical coordination and management for the Trinity contract
- Deliver and install Trinity
- Continue Trinity development and engineering (D&E) collaborations between selected vendor and ACES

FY16

- Complete Trinity System Integration Readiness Level 2 milestone
- Provide technical coordination and management for the Trinity contract
- Enable Trinity operations in support of Advanced Technology Computing Campaigns (ATCCs)
- Continue Trinity R&D collaborations between selected vendor and tri-labs

FY17

- Run two ATCCs
- Provide technical coordination and management for the Trinity contract
- Complete Trinity Production Readiness Level 2 milestone
- Submit Trinity CD4 to the DOE and obtain approval
- Continue Trinity operations in support of ATCCs

FY18

- Run two ATCCs
- Provide technical coordination and management for the Trinity contract
- Continue Trinity operations in support of ATCCs

FY19

- Run two ATCCs
- Provide technical coordination and management for the Trinity contract
- Continue Trinity operations in support of ATCCs

Alliance for Computing at Extreme Scale Cielo Capability Computing Platform (LANL, SNL)

The Cielo capability computing platform is a project under the ACES. ACES is a joint collaboration between LANL and SNL defined under an MOU to provide a user facility for capability computing to the NNSA weapons programs in support of stockpile stewardship, to develop requirements and system architecture for ASC capability systems requirements definition, architecture design, procurement, key technology development, systems deployment, operations, and user support.

The architecture and design of Cielo is optimized to provide performance at the full scale of the machine, in support of the NNSA program's most challenging CCCs. This project covers all aspects of the technical, programmatic, and procurement planning for the platform.

Cielo is the primary platform that supports the ASC CCC. In FY10 it replaced the Purple platform at LLNL and is sited at LANL and operated by ACES. Cielo provides 1.37 peak petaFLOP/s with over 140,000 compute cores, and 10 petabytes of storage. Over 6,000 of the cores are dedicated to visualization and data services activities with connections to the SNL and LLNL sites.

Required Capabilities

R1: Architecture, design, and operational support under the ACES MOU

R2: Production, capability-class computing support for the ASC's Capability Computing Campaigns

Five-Year Plan

FY15

- Complete the Cielo campaigns for FY15 (CCC7 and CCC8)
- Continue to run Cielo in production capability mode
- Provide operations in support of CCCs

FY16

- Provide operations in support of CCCs
- Retire Cielo after successful deployment of Trinity

FY17

None. Project completed.

Next-Generation Computing Technologies (WBS 1.5.4.9)

The Next-Generation Computing Technologies product includes costs for the planning, coordinating, and executing of the next-generation R&D computing technology activities. These activities will prepare the ASC applications and computing environment for the next computing paradigm shift to extreme parallelism, via heterogeneous and/or multi-core nodes.

Next-Generation Computing Environment (LLNL)

The Next Generation Computing Enablement efforts will help prepare ASC for the post-petascale era, addressing the software environment and platforms. It includes gathering requirements for post-petascale computing and coordinating next generation activities internally and externally. The efforts will enable advanced application work to develop benchmarks for new platforms as well as to adapt codes to the expected new architectures. The software efforts are focused on an architecture that ties together system level software, resource management, development tools, data analysis tools, and programming models, while addressing ASC application requirements. On the hardware side, these efforts include tracking and collaborating on technology innovations. This effort includes interactions with ASCR, vendors, and academia, including planning and technical coordination for vendor contracts. Team members will carry out investigations and co-design activities using test beds and existing technology, making use of proxy applications.

Required Capabilities

R1: Detailed description of post-petascale requirements across software and hardware environment

R2: Understanding challenges the ASC program must overcome to field next-generation systems

R3: Prototypes and simulators to assess the impact on ASC applications

Five-Year Plan

FY15

- Develop next-generation CTS and run-time level support, including center-wide resource management, accelerator support, and performance monitoring
- Provide technical coordination for 2017 ASC ATS contracts
- Identify technology gaps in available systems, and procure and deploy prototypes and test beds to allow exploration of how those technologies support ASC applications
- Refine and begin execution of LLNL plan for the software environment for the Sierra ATS

FY16

- Conduct co-design activities with vendors, and research and evaluate next-generation technologies
- Provide technical coordination for 2017 ASC ATS contracts
- Coordinate and administer FastForward2 contracts and participate in DesignForward activities

- Identify technology gaps in available systems and procure and deploy prototypes and test beds to allow exploration of how those technologies support ASC applications

FY17

- Conduct explorations on 2017 platform
- Conduct co-design activities with vendors, and research and evaluate next-generation technologies
- Prepare to site 2017 ASC ATS
- Coordinate and administer contracts, such as FastForward2
- Identify technology gaps in available systems and procure and deploy prototypes and test beds to allow exploration of how those technologies support ASC applications

FY18

- Evaluate next-generation technologies as delivered on the ASC ATS
- Identify gaps and enhancements needed for technologies in co-design effort with ASC IC
- Coordinate and administer contracts that lay the groundwork for the next-generation machines beyond the Sierra ATS
- Identify technology gaps in available systems and procure and deploy prototypes and test beds to allow exploration of how those technologies support ASC applications

FY19

- Explore advanced technology developments
- Administer contracts for next-generation research in software environments
- Conduct co-design activities with IC, other labs, and vendors
- Exercise applications on test beds to identify gaps and enhancements needed for next-generation IC codes

Future Architecture Planning and System Requirements (LANL)

The major focus of the Future Architecture Planning and System Requirements project is to define requirements and potential system architectures for advanced systems platforms that meet ASC programmatic requirements and drivers. This project covers all aspects of program and procurement planning for current and advanced systems and strategic planning for supporting infrastructure. Additionally, this project provides a focus for the various planning efforts. In FY15, this project will focus on the project management of the ASC system called Trinity. The focus in this project also includes the execution of DOE Order 413.3.

Required Capabilities

R1: Working with vendors and users, identify potential characteristics of platform architectures to meet mission need

Five-Year Plan

FY15

- Provide program and project management for computing platforms, including requirements gathering and analysis
- Plan infrastructure to support pre-exascale and exascale systems
- Participate in site-wide planning for power and cooling upgrades for future systems

FY16

- Provide program and project management for computing platforms, including requirements gathering and analysis
- Continue planning for ATS-3 platform

FY17

- Provide program and project management for computing platforms, including requirements gathering and analysis

FY18

- Provide program and project management for computing platforms, including requirements gathering and analysis

FY19

- Begin the ATS-3 acquisition process

Future-Generation Computing Technologies (LANL)

This project includes high-risk, high-reward research for future systems, including research on advanced programming models and hierarchical storage technology for advanced in-situ, and support of DOE initiatives such as FastForward and DesignForward.

Next-Generation Computing Technology is aimed at high-risk, high-reward investigations that can enable ASC codes on new system architectures. This includes interaction with vendors on FastForward, DesignForward, and other such initiatives. The objective of the PINION project is to investigate the use of high-level data parallelism in the implementation of physics algorithms of interest to the ASC Program. With this model, algorithms are written using a relatively small set of data-parallel primitive operators, such as transform, reduce, and scan, along with custom functors. Backend implementations of the data-parallel primitives optimized for specific architectures then allow these higher level physics codes to be portable across these architectures, making efficient use of multi-core and many-core parallelism available on each. The investigations undertaken by the PINION project are expected to have high relevance for the future of ASC codes. Hierarchical storage technology for advanced analytics will investigate alternative hierarchical storage technologies, such as novel burst-buffer designs for support of in-situ and in-transit analysis and streaming analytics. Interactions with vendors for FastForward and DesignForward are included, such as gathering of computation and communication traces and support of mini-apps.

Required Capabilities

R1: Provide novel programming models to expose the massive parallelism of new architectures, explore focused functional capabilities of the ASC code base (for example, unstructured meshes and particle operations)

R2: Provide co-design knowledge base consisting of recommendations and best practices for code development on Trinity and future extreme-scale architectures

R3: Working with vendors, explore areas where custom hardware can be applied to applications of interest and develop prototypes for demonstration

R4: Design and demonstration of low-level runtime components, for example, lightweight-thread scheduling and memory affinity on both shared and distributed memory architectures

R5: Language and compiler design and development (including memory and process/thread management), focused on needs of IC workloads

Five-Year Plan

FY15

- Investigate leveraging cloud and big data technology for ASC data analytics
- Research NVRAM for memory and storage hierarchy

- Investigate alternative burst buffer designs for in-situ
- Develop data structures and algorithms for physics code and backend optimizations for data and task parallel

FY16

- Prototype multiple algorithms of interest to ASC in THRUST and PINION
- Explore distributed data systems to support streaming analytics

FY17

- Work with IC to validate the programming model and show performance improvements to algorithms of interest to ASC
- Continue to explore advantages of distributed data systems for HPC
- Continue interaction with vendors on advanced architectures

FY18

- Work with IC to validate the programming model and show performance improvements to algorithms of interest to ASC; extend into IC codes
- Demonstrate advantages of distributed data systems for HPC
- Continue interaction with vendors on advanced architectures

FY19

- Work with IC to validate the programming model and show performance improvements to algorithms of interest to ASC
- Demonstrate advantages of distributed data systems for HPC
- Continue interaction with vendors on advanced architectures

Architecture Office (SNL)

The objective of this project is to analyze potential computer and system architectures for platforms that meet future ASC programmatic requirements for ATS-3 and beyond. The primary activity is to establish a technology foundation for ASC to influence the directions for future hardware and system software architectures for ASC ATS. The project will track HPC computer industry hardware/software trends with a specific focus on the identification of opportunities to influence future hardware architectures and development of future system software that provides an on-ramp for the ASC application code base. This project is also the focal point for the active collaboration of SNL technical staff with industry FastForward, DesignForward, and potential future PathForward R&D projects.

Required Capabilities

R1: Translation of mission and application drivers into architecture and system software requirements

R2: Support for SNL CSSE, IC, Physics and Engineering Models (PEM), ASC co-design project, and ATDM staff to collaborate with industry hardware architecture R&D projects, for example, FastForward, DesignForward, and other industry R&D projects

R3: Understanding of industry roadmaps and opportunities to influence technologies available in the ATS-3 timeframe and beyond

Five-Year Plan

FY15

- Support SNL engagement with industry *Forward R&D projects
- Collaborate technically and coordinate with the joint ASC/ASCR R&D project with Micron Technology
- Initiate discussions with industry to determine both technology roadmaps and opportunities for influencing future hardware architectures for ATS-3 and beyond

FY16

- Continue to support SNL engagement with industry *Forward R&D projects
- Collaborate with Micron on Cube Module and Multi-Cube Functions in memory design studies
- Document architecture analysis for potential ATS-3 alternatives and opportunities to influence hardware architectures for ATS-4 and beyond

FY17

- Continue to support SNL engagement with industry *Forward R&D projects

- Support DOE ASC/ASCR request for proposal (RFP) for future hardware architecture R&D projects

FY18

- Continue to support SNL engagement with industry *Forward R&D projects
- Initiate discussions with industry to determine both technology roadmaps and opportunities for influencing future hardware architectures for ATS-5 and beyond

FY19

- Continue to support SNL engagement with industry *Forward R&D projects
- Document architecture analysis for potential ATS-5 alternatives
- Analyze and document recommendations for Beyond CMOS technology opportunities

Heterogeneous Computing (SNL)

The Heterogeneous Computing project will develop capabilities that facilitate ASC applications' ability to take advantage of heterogeneous architectures with many-core-accelerators, including NVIDIA graphics processing unit (GPU), Intel many-integrated core (MIC), and multicore CPUs. There are clear challenges for coupled, multi-physics-based simulation incorporating unstructured meshes and implicit solution methods. This project will work in tight alignment with co-design efforts and proxy-application development to explore performance on available heterogeneous architecture test beds. Programming models and associated runtime support for portable hybrid parallelism and data locality/placement must be developed to efficiently exploit the diverse set of many-core processors proposed for these architectures. In addition, fine-grained, architecture-aware load balancing of work must be explored and utilized.

These multi-physics coupled applications have a higher level of heterogeneous parallelism that is not currently being exploited. This project will explore SNL's ability to support parallel invocation of coupled applications and efficient management of the shared data. Both physics and engineering applications critical to the nuclear weapons enterprise must be able to take advantage of all levels of parallelism to benefit in the context of exascale computing.

This project will build upon researchers' expertise to improve application performance and portability to the next generation of architectures for scaling applications to a billion-way parallelism. This project has close ties to the heterogeneous architectures test beds in SNL's Advanced System Technology R&D project.

Required Capabilities

R1: A performance-portable programming model that enables a large body of complex, fine grained parallel computational kernels to be efficiently ported to the variety of emerging many-core accelerators, and obtain good performance on these devices with no or negligible code modifications

R2: Proxy-applications that fully exercise this programming model to enable early evaluation of new generations of many-core accelerators

Five-Year Plan

FY15

- Collaborate with library and application teams to prototype strategies to migrate their capabilities to heterogeneous computing architectures
- Research hierarchical, heterogeneous domain decomposition to address Non-Uniform Memory Access (NUMA) performance concerns and work decomposition in an MPI + central processing unit (CPU) + accelerator strategy

FY16

- Research deeper, heterogeneous parallel work hierarchies for complex multi-physics applications; research programming models to manage the complexity in such hierarchies
- Develop proxy-applications to demonstrate fully utilizing all levels of heterogeneous parallel capabilities within a compute node

FY17

- Develop proxy-applications with deep heterogeneous parallel work hierarchies

FY18

- Develop proxy-apps to evaluate programming models for processor-in-memory architectures

FY19

- Analyze and report the performance results from running proxy applications on emerging test beds with deep parallel work hierarchies and different memory architectures.

System Software and Tools (WBS 1.5.4.4)

This level 4 product provides the system software infrastructure, including the supporting OS environments and the integrated tools, to enable the development, optimization, and efficient execution of application codes. The scope of this product includes planning, research, development, integration and initial deployment, continuing product support, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include system-level software addressing optimal delivery of system resources to end-users, such as schedulers, custom device drivers, resource allocation, optimized kernels, system management tools, compilers, debuggers, performance tuning tools, run-time libraries, math libraries, component frameworks, other emerging programming paradigms of importance to scientific code development and application performance analysis.

System Software Environment for Scalable Systems (LLNL)

The System Software Environment for Scalable Systems project provides system software components for all the major platforms at LLNL, research and planning for new systems and future environments, and collaborations with external sources such as the platform partners, especially IBM and Linux vendors. This project covers system software components needed to augment Linux and required proprietary operating systems that function in a manageable, secure, and scalable fashion needed for LLNL ASC platforms.

This project includes work on developing, modifying, and packaging the Tripod Operating System Software (TOSS), and developing scalable system management tools to support the OS and interconnect (for example, TOSS and IB monitoring tools), as well as the resource management environment (Moab and Simple Linux Utility for Resource Management (SLURM)) to queue and schedule code runs across LLNL systems. LLNL uses TOSS on all of its Linux clusters. This project also funds approximately 60 percent of the manpower required to develop, deploy, and maintain TOSS. The funding LLNL receives for its portion of Facility Operations and User Support's (FOUS') TOSS funding accounts for 40 percent of the effort required to develop, deploy, and maintain TOSS. Therefore, TOSS activities and deliverables at LLNL are captured both here and in section 1.5.5.6 of this document.

Required Capabilities

R1: Fully functional cluster software (kernel, Linux distribution, IB stack and related libraries, and resource manager, and cluster-management tools)

R2: Capable of running MPI jobs at scale on Linux capacity clusters

R3: Full lifecycle support, including release management, packaging, quality assurance (QA) testing, configuration management, and bug tracking

Five-Year Plan

FY15

- Provide ongoing TOSS software development and support
- Develop/deploy TOSS 2.X (based on RHEL 6.X)
- Develop/deploy TOSS 3.X (based on RHEL 7.X)
- Develop enhancements to the IB OpenSM Monitoring Service, which provides an application interface to IB diagnostics, monitoring, management, and control functions
- Investigate alternative architectures for commodity Linux clusters (for example, ARM)

FY16

- Provide ongoing TOSS software development and support

- Develop/deploy TOSS 3.X (based on RHEL 7.X)
- Initiate development of TOSS 4 (based on RHEL 8)
- Develop/deploy identified system software projects for efficient operation ~10,000 node scale

FY17

- Provide ongoing TOSS software development and support
- Develop/deploy TOSS 4.X (based on RHEL 8.X)

FY18

- Provide ongoing TOSS software development and support
- Initiate development of TOSS 5 (based on RHEL 9)

FY19

- Provide ongoing TOSS software development and support
- Develop/deploy TOSS 5.X (based on RHEL 9.X)

Applications Development Environment and Performance Team (LLNL)

The Applications Development Environment and Performance Team (ADEPT) project provides the code development environment for all major LLNL platforms, supports user productivity, provides research and planning for new tools and future systems, and collaborates with external sources of code development tools. The project works directly with code developers to apply tools to understand and to improve code performance and correctness. The elements of the development environment covered by this project include, but are not limited to, compilers, debuggers, performance and memory tools, interfaces to the parallel environment, and associated run time library work.

Required Capabilities

R1: Robust development environment to support ASC applications running on all major platforms

R2: Tools for researching application code performance and correctness on current and future systems

R3: Compilers and runtimes for all major programming languages and programming models, including C, C++, Fortran, MPI, and OpenMP

R4: Software infrastructure to implement tools and runtimes

R5: Production-level support for checkpoint/restart and parallel I/O

Five-Year Plan

FY15

- Begin scoping of 2017 system requirements for development environment
- Research and further develop resilience strategies based on SCR Lib for 2015 and future systems
- Provide support for identified emerging programming models
- Plan for development environment for CTS-1
- Support tri-lab code teams in CCC activities on Sequoia

FY16

- Work with vendor partner identified for 2017 system in co-design for the development environment for the platform
- Research and develop correctness tools for emerging technologies
- Support CTS-1 environment

FY17

- After initial system delivery, deploy development environment, test, and validate with user codes

- Continue development of tools for correctness, resilience, performance as gaps are identified
- Support CTS-1 environment

FY18

- Support user codes on porting, debugging, and performance tuning efforts on the delivered Sierra ATS
- Support CTS-1 environment
- Support the Sierra ATS environment

FY19

- Support user codes on performance tuning efforts on Sierra and CTS-1 platforms
- Provide support for emerging programming models for ATS platforms
- Plan for development environment for CTS-2

High Performance Computing Systems Research (LANL)

HPC systems research is a broad-reaching project focusing on near to long-term research of all the components needed to support a rich environment for very large-scale applications. Systems research bridges the gap between hardware and programming model, and requires tight collaboration in supporting the development of programming models, tools, visualization/analytics, and system software aspect of I/O.

The project includes investigations on resilient system services, soft-error resilience, system support for data-intensive computing, power, and interconnect topology modeling/evaluation.

Resilient system services focus on developing a vehicle to investigate resilient, dynamic, distributed, scalable services for large-scale systems and providing an interface to programming models so that ASC applications can access these features on current and future hardware. Current activities include investigation of distributed systems software for job launch and monitoring. An imminent challenge for extreme-scale HPC systems is the issue of power limits and rapidly varying demands on the grid. Techniques for power-capping HPC systems will be investigated.

Investigations of soft-error resilience will continue work on an accelerated testing environment for soft error profiling using a virtual machine (VM) approach to inject faults while actively running real ASC codes. Hardware reliability in HPC systems remains a challenge to characterize. Statistical studies of reliability data from a variety of production systems will be extended, and models of dynamic random access memory (DRAM) reliability will be developed. Reliability of non-volatile storage will be studied in support of future hierarchical storage systems (for example, burst buffer architectures).

Required Capabilities

R1: Distributed systems software expertise

R2: OS support for future HPC applications and runtimes

R3: Resilience of software and systems

R4: Expertise in power and interconnect issues

Five-Year Plan

FY15

- Analyze and develop advanced HPC system software, including persistence of deep memory hierarchies, modeling interconnects, and mixed scheduling with traditional MPI and many-task (fine grained-ensemble) workloads
- Analyze and extend models of DOE production system statistics related to reliability, including system characterization, code characterization, and SRAM aging
- Implement and extend data-parallel methods via extensions to Nvidia's Thrust library using hydrodynamic algorithms using AMR meshes as a representative algorithm

- Develop in-transit analysis workloads and a data-aware computing interface to allow ASC codes to leverage data-intensive compute models

FY16

- Analyze and develop advanced HPC system software, including persistence of deep memory hierarchies and NVRAM
- Analyze and extend models of DOE production system statistics related to reliability including: system characterization, code characterization, and static random access memory (SRAM) aging
- Model, simulate, and evaluate topology

FY17

- Investigate integration of node OS and system services into cohesive distributed system
- Continue to characterize and model reliability of emerging technologies for HPC
- Provide OS and runtime support for emerging programming models for next-generation systems

FY18

- Improve performance and scalability of global OS and system services
- Provide OS and runtime support for emerging programming models for next-generation systems
- Continue to characterize and model reliability of emerging technologies for HPC

FY19

- Improve performance and scalability of global OS and system services
- Provide OS and runtime support for emerging programming models for next-generation systems
- Continue to characterize and model reliability of emerging technologies for HPC

Advanced System Test Beds (LANL)

The Test Beds project provides program management for CSSE and provides test bed hardware and software for research investigations in support of the CSSE mission. It fills the gaps of advanced architecture hardware and provides local access to advanced hardware.

Required Capabilities

None

Five-Year Plan

FY15

- Manage advanced technology computer test beds for CSSE and IC project use
- Refresh technology in Darwin test bed system with state-of-the-art technology and vendor loans of future products, focusing on representing most likely hardware for ATS-1, ATS-2, and ATS-3

FY16

- Provide local access to advanced architectures
- Refresh hardware in test bed and evaluate tri-lab to fill gaps

FY17

- Provide local access to advanced architectures
- Refresh hardware in test bed and evaluate tri-lab to fill gaps

FY18

- Provide local access to advanced architectures
- Refresh hardware in test bed and evaluate tri-lab to fill gaps

FY19

- Provide local access to advanced architectures
- Refresh hardware in test bed and evaluate tri-lab to fill gaps

System Software Stack Advancement (SNL)

The System Software Stack Advancement project supports system software R&D to address scalability and efficiency of future computational systems in multiple dimensions. An important aspect is providing lightweight services and functionality that does not compromise scalability and therefore performance. The focus will be on three critical areas for HPC systems, which will enhance efficiency, performance and scalability of applications on future HPC systems:

- Power has become a first-order design constraint for future supercomputers. SNL will expand upon work in data collection and tuning techniques that provided new insight into understanding power requirements and affecting power use of ASC applications.
- Much of the complexity of managing future heterogeneous compute and memory resources will the responsibility of the runtime system. SNL will continue to explore the relationship between the runtime system, the operating system, and the interconnect to provide the necessary policies and mechanisms for ensuring scalability and performance while insulating the complexities of the resources from applications.
- Previous work with virtualization has shown promise in the area of HPC. Virtualization will be leveraged to provide insights into application runtime characteristics and where optimization efforts would be best targeted.

As a long-term goal, SNL plans to integrate these targeted efforts with previous successes in lightweight operating systems (Kitten), lightweight runtime systems (Othreads), and high performance network stack (Portals communication protocol) development with a production HPC computing stack. While this is a significant development effort, the long-term benefits are many:

- Risk mitigation against vendor provided software failure
- Proof-of-concept demonstration of scalable or efficiency enhancements at scale
- Setting the standard for vendor-delivered scalable software stacks
- Enabling scalable research in both systems software and other areas dependent on scalable software features
- More complete understanding of the responsibilities of system software relative to applications
- Development of the necessary abstractions, interfaces, and mechanisms needed between the application and the system software as well as the system software and the underlying hardware

This effort will necessarily be accomplished in conjunction with the acquisition of future ATS.

Required Capabilities

R1: New software features and capabilities are required to maintain and enhance scalability, productivity, and efficiency of applications on advanced architecture platforms

Five-Year Plan

FY15

- Perform benchmark analysis of node virtualization layer running on a production ASC platform
- Integrate power/energy measurement and API on Trinity's early-spec platform
- Evaluate effectiveness of LWK and dynamic adaptive runtime system combination for ASC workloads
- Prototype portions of the published Power API specification using PowerInsight as the measurement and control mechanism
- Evaluate the integration of a dynamic runtime system with interconnect capabilities

FY16

- Demonstrate application portions of power/energy measurement and API on ACES Trinity platform
- Characterize and demonstrate increased application energy efficiency on Trinity
- Port scalable virtual machine environment to Trinity
- Demonstrate improved application performance from integration of a dynamic runtime system with compute, memory, and network resources

FY17

- Perform benchmark analysis of node virtualization layer on Trinity
- Deploy production capability of power/energy measurement and API on ACES Trinity
- Develop abstractions for applications to communicate policies for application adaptivity to a dynamic runtime system

FY18

- Evaluate LWK with production workloads at large scale on Trinity platform
- Demonstrate an initial implementation of the interface for communicating application policies for dynamic runtime system adaptivity

FY19

- Demonstrate an advanced application workflow using LWK with virtualization

- Evaluate application performance using the advanced interfaces for runtime system adaptivity

High Performance Computing Hardware Architecture Simulation (SNL)

The SST is a suite of tools enabling multi-scale computer architecture simulation to meet the needs of HPC software/hardware co-design. The SST consists of a core set of components that enable parallel discrete-event simulation; high-fidelity networking, memory, and processor components; and coarse-grained simulation components that capture essential elements of machine performance with low computational cost. Future HPC systems and the applications designed to utilize them are impacted by a variety of considerations, including scalability of applications, ease-of-programming, memory and network latencies becoming more imbalanced relative to computation rates, data corruption and its propagation, frequency of interrupts, power consumption, and overall machine cost. SST is designed to allow each of these parameters to be explored, permitting the consideration of a broad space of potential architectural and application/algorithmic designs. The goal is for the SST components to be extended and enhanced by a community of simulator developers, including academic, industrial, and government partners. An even larger community is expected to be the users of SST, including algorithm developers, architecture designers, and procurement team members.

Required Capabilities

R1: Ability to analyze and understand potential performance of large-scale HPC systems before they are built

R2: Ability to support proactive and holistic co-design and provide prioritized recommendations for future hardware and system architecture capabilities that would have the greatest impact on ASC applications and mission drivers

Five-Year Plan

FY15

- Report on architectural analysis of Trinity components (pre-test)
- Provide a pre-configured, "abstract machine model" and associated proxy architecture parameters that proxy application developers can use for co-design analysis
- Merge SST (macro/micro) with a unified, parallel discrete event simulator core and a common component interface to allow integration of macro or micro components into the HPC architectural simulation framework

FY16

- Demonstrate a use case for unified SST architectural simulation to analyze the Trinity system and compare with Trinity measurements with selected proxy applications
- Create an SST model that can aid in the analysis of reliability of for example, processor components and memory buses
- Provide a power and energy consumption analysis capability

FY17

- Down select hetero-core models to the one or two that have gained traction in the recent years
- Demonstrate an SST simulation of a proxy application that was migrated from an existing form (MPI-everywhere) to an exascale-compatible programming/execution model
- Report on comparison of SST Hybrid Memory Cube (HMC) analysis with HMC test-bed measurements

FY18

- Using SST, analyze and report on the tradeoffs of heterogeneous big-core and little-core processors
- Demonstrate SST release with Dakota UQ capability (for selected SST components)

FY19

- Use SST to model and analyze post-CMOS components integrated into a 2022-era advanced architecture concept

Interprocess Communication System Software Stack (SNL)

The Interprocess Communication System Software Stack project will develop capabilities to enable performance and scalability of ASC applications on current and future high-performance interconnection networks on extreme-scale platforms. This project will concentrate on characterizing application requirements with respect to functionality and performance for intra-application data movement as well as application network transfers to external I/O services. It will also provide a low-level network programming interface appropriate for current-generation network hardware as well as more advanced next-generation hardware with more sophisticated network interface capabilities and functionality. As applications explore alternative programming models beyond the current distributed memory MPI model, the low-level network programming interface must evolve to include the ability to provide very lightweight one-sided data transfer operations, while continuing to enable efficient two-sided message-based transfers. It is likely that this project will expand to include an analysis of network topologies, network interface hardware design and evaluation, optimized network transfer protocols, and system software support for advanced network interface operations.

This project will build on existing efforts surrounding the development of the next-generation Portals network programming interface and measurements of application sensitivity to network performance. It has close ties to the ACES project and the Cray network hardware design and engineering activity in WBS 1.5.4.1.

Required Capabilities

R1: Future interconnects for extreme-scale platforms must enable performance and scalability for ASC applications

R2: A low-level interconnect programming interface that meets the functionality, performance, and scalability requirements of applications and system services

Five-Year Plan

FY15

- Continue to develop and enhance Portals to meet the ongoing interconnect requirements of applications and services
- Study the impact of low-power processors on interconnect performance
- Explore power/energy tradeoffs in the communication software stack

FY16

- Provide a set of interconnect performance requirements for important ASC applications
- Enhance the Portals specification to better support simultaneous access to the interconnect from tens to hundreds of threads

FY17

- Demonstrate a hardware implementation of Portals
- Deploy a reference implementation of the enhanced Portals specification for supporting tens to hundreds of threads at network endpoint

FY18

- Complete a study that evaluates the performance of the hardware implementation of Portals
- Deploy a prototype implementation of MPI endpoints using the enhanced Portals specification

FY19

- Complete a study that examines the key performance and scalability issues with hundreds of applications threads simultaneously accessing interconnect resources

Resilience (SNL)

The next generation of computing platforms promises both new capabilities and increased capacity for meeting SNL's mission challenges. However, these platforms will involve new computer architectures. It is expected that the reliability of these systems may be degraded by both the shear number of components as well as their susceptibility to errors as feature sizes are pushed to the limit. This project explores possible solutions to provide resilience to system errors that will enable our new ATDM codes to effectively use the new computational hardware.

Note: this description reflects early planning for ATDM and we expect that required capabilities and detailed deliverables will be refined and extended as ATDM planning continues to mature.

Required Capabilities

R1: Advanced, flexible check pointing strategies that can utilize new hardware features (such as burst buffers)

R2: Algorithms that are fundamentally resilient to hardware failures

R3: Accurate methods to model and access the impact of resilience mitigation mechanisms on application performance

R4: The ability to define persistent data object stores that will support scalable fault management at the application level

R5: The ability for application programmers to declare specific segments of code to execute reliably with a designated level of assurance

Five-Year Plan

FY15

- Demonstrate research code framework incorporating one or more representative candidates for robust PDE solvers and hardware silent-error models expected to be applicable at extreme scale
- Develop proof-of-concept example for using symbolically represented constraints to automatically generate a solver algorithm with resilience to a hardware silent-error model
- Deploy local failure/local recovery technologies for checkpoint/restart in a mini-driver application to investigate and qualify its impact on application architecture/design
- Assess the LFLR model with more complex simulation code running on the production systems along with more realistic hard failures, including a loss of multiple processes and alternate recovery (for example, roll-forward) schemes
- Develop and demonstrate an application performance model based on operating system interference to evaluate the performance of resilience mitigation mechanisms

- Using the performance model, characterize the CPU detour footprint of LFLR and communication characteristics most impacted by this mitigation method

FY16

- Develop needed parameterization of hardware silent-error model to systematically transform algorithms into resilient numerical implementations
- Implement LFLR persistent storage API on leadership-class system
- Define prototype selective unreliability API and semantics
- Demonstrate vertically integrated application resilience, including prototype reliable code segment specification and integrated fault models and detectors with system response mechanisms
- Extend and demonstrate the CPU detour performance model for additional failure mitigation methods (for example, selective reliability, task rescheduling, and load balancing)

FY17

- Demonstrate a representative robust simulation application at scale on an advanced technology system and evaluate its performance and silent-error tolerance
- Impact co-design of next-generation HPC architectures to optimize support for efficient resilience mechanisms
- Specify APIs, requirements, and execution model for integrated resilience capabilities, initial production versions of core capabilities
- Demonstrate vertically integrated application resilience on a production system of an ASC IC application or library component
- Extend the performance model to account for impact of next-generation HPC architecture capabilities

FY18

- Deploy integrated resilience capabilities in production for ASC codes on extreme-scale systems

FY19

- Demonstrate an intrinsically robust adaptive HPC software stack that "programs itself," optimizing performance while maintaining specified correctness properties on unreliable hardware

Input/Output, Storage Systems, and Networking (WBS 1.5.4.5)

This level 4 product provides I/O (data transfer) storage infrastructure in balance with all platforms and consistent with integrated system architecture plans. The procurement of all supporting subsystems, data transfer, storage systems, and infrastructures occurs through this product. The scope of this product includes planning, research, development, procurement, hardware maintenance, integration and deployment, continuing product support, quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include high-performance parallel file systems, hierarchical storage management systems, storage-area-networks, network-attached storage (NAS), and high-performance storage system (HPSS) or future hierarchical storage management system disks, tape, robotics, servers, and media. This product also includes relevant prototype deployment and test bed activities. Projects and technologies in the advanced networking and interconnect areas include networking and interconnect architectures, emerging networking hardware technologies and communication protocols, network performance/security monitoring/analysis tools, and high performance encryption and security technologies.

Archive Storage (LLNL)

The Archival Storage project provides long-term, high-performance, archival storage services to ASC customers. This includes a collaborative software development effort (currently, HPSS) between the tri-labs, ORNL, Lawrence Berkeley National Laboratory, and IBM, as well as deployment and support of archival storage software and interfaces for tri-lab ASC customers on unclassified and classified networks. It includes the selection, procurement, deployment, support, and maintenance of archival storage hardware and media, ongoing technology refresh, and data stewardship. HPSS provides scalable, parallel, archival storage interfaces and services to the tri-labs.[2]

A world-class array of hardware is integrated beneath HPSS, supplying the performance necessary to offload ASC platforms, thereby increasing computation. This includes disk arrays, robotic tape subsystems, servers, storage area networks (SANs), networks, and petabytes of tape media, enabling high-speed parallel transfers into a virtually unlimited data store.

Required Capabilities

R1: Cost-effective, long-term, high-performance archival storage services for ASC customers

Five-Year Plan

FY15

- Continue ongoing HPSS software development and support, with focus on inspection, integration testing, and release of HPSS 7.5, which features partitioned metadata
- Begin planning for production deployment of HPSS 7.5, including evaluation and potential procurement and deployment of new HPSS Core Server platforms and metadata disk subsystems to exploit scalability of partitioned metadata
- Develop and deploy hard quota and user notification features in archival quota system
- Begin repack of nine-year-old T10K Gen1 media (1-TB native) to T10K Gen2 media (8-TB native) to minimize data loss due to aging media and to reclaim slot capacity in libraries
- Deploy HPSS disk subsystems to increase disk cache capacity and disk file residency
- Deploy network upgrades from 10Gbe to 40Gbe for increased archive performance
- Provide ongoing support of currently deployed archival storage systems, including selection, deployment, support, and maintenance of all archival storage hardware and media, customer and interface support, ongoing tech refresh, and data stewardship

[2] See http://www.hpss-collaboration.org/index.shtml.

FY16

- Finalize, test, and release future version of HPSS featuring checksums and distributed servers and metadata for increased scalability of archive metadata performance
- Evaluate, select, procure, and deploy hardware upgrades to HPSS core server, metadata, and disk subsystems for increased archive metadata performance, capacity, and bandwidth

FY17

- Develop future version of HPSS, including additional metadata scalability features in preparation for exascale computing needs
- Deploy future version of HPSS featuring checksums and distributed servers and metadata for increased scalability of archive metadata performance
- Evaluate, select, procure, and deploy upgrades to enterprise tape drive environment for increased archive capacity and bandwidth

FY18

- Finalize, test, and release future version of HPSS featuring additional metadata scalability features in preparation for exascale computing needs
- Deploy network upgrades from 40Gbe to 100Gbe for increased archive performance

FY19

- Deploy future version of HPSS, including additional metadata scalability features in preparation for exascale computing needs
- Evaluate, select, procure, and deploy hardware upgrades to mid-range tape drive environment for increased archive capacity and bandwidth

Parallel and Network File Systems (LLNL)

The Parallel and Network File Systems (NFS) project provides for the development, testing (feature, capability, performance, and acceptance), procurement, integration, and ongoing support of various file system technologies and interfaces necessary for the efficient and effective use of ASC high-performance platforms. Included are the continuing development and support of Lustre as a fully featured file system for the range of ASC platforms, the deployment and support of ubiquitous NAS services for home, project, and scratch space, and the I/O support of various programming interfaces for parallel I/O.

This project deploys and supports Lustre file systems for ASC platforms as well as high-availability NAS file systems for home and project space, and scratch space for serial capacity clusters. It actively works with the OpenSFS Lustre development community to add Lustre file system scalability and reliability enhancements required by ASC platforms. The file system up through the programming interfaces are supported to help developers of applications use parallel I/O effectively.

Required Capabilities

R1: Provide reliable and available high-speed parallel file system access in support of ASC platforms

R2: Maintain optimal Lustre and NFS functionality across Center machines for local and remote users

R3: Develop and deploy Lustre HPC-focused software releases in concert with broader Lustre community

R4: Deploy new file systems consistent with platform requirements

R5: Work with users and application developers on developing and improving I/O access mechanisms

Five-Year Plan

FY15

- Enhance ZFS-based Lustre metadata performance in support of user and purge performance
- Support the development, testing, and deployment of new Lustre versions in classified and unclassified environments
- Complete analysis of reliability and availability realized through the selective mounting of file systems on SCF platforms
- Deploy file system resources in support of CTS-1 deliveries
- Complete deployment of new NAS home directory hardware in open computing facility (OCF) and secure computing facility (SCF) centers
- Investigate pNFS production viability leveraging the Hyperion test environment

FY16

- Continue HPC-focused development in concert with file system community
- Field new file system(s) in support of ASC platform requirements
- Prepare file system solution for Sierra ATS
- Support production implementation of tri-lab shared file system

FY17

- Continue HPC-focused development in concert with file system community
- Field new file system(s) in support of ASC platform requirements
- Deploy file system for Sierra ATS

FY18

- Continue HPC-focused development in concert with file system community
- Field new file system(s) in support of ASC platform requirements

FY19

- Continue HPC-focused development in concert with file system community
- Field new file system(s) in support of ASC platform requirements

Networking and Test Beds (LLNL)

The Networking and Test Beds project provides research, performance testing, capability testing, and analysis for the file system, network, and interconnect subsystems in support of current and future systems and environments. This work relies heavily on an adequately provisioned test bed, skilled staff, and collaborations with vendors.

This project will test various hardware and software components to quantify the features, performance, reliability, security, and interoperability of the products and broader technology base. The information acquired as a result of this project will be used to help determine an integrated architecture and resultant procurements for these subsystems.

Required Capabilities

R1: Adequately provisioned test bed for researchers and developers to perform testing and analysis in support of current and future systems and environments

Five-Year Plan

FY15

- Deploy small ARM64 cluster for software development and debugging working with RedHat; continue to address software shortcomings
- Support Intel Haswell in TOSS 2 and 3; test and evaluate
- Integrate RHEL7 into TOSS 3 and test
- Evaluate Ares, Stormlake, and Mellanox enhanced data rate (EDR)
- Evaluate and support TOSS 3 in CTS-1 systems

FY16

- Evaluate new processor technology

FY17

- Evaluate new memory architectures, if available

FY18

- Evaluate new processor technology; evaluate creating HPC-specific CPU based on licensed IP

FY19

- Test RHEL8
- Work with RedHat vendor to test new CPU and network technologies

File Systems, Archival Storage, and Networking (LANL)

Capabilities of the Archival and File Systems components of the project include online file systems such as the NFS complex and enterprise-wide supercomputer file systems, general parallel file system (GPFS) development, deployment and management, scalable I/O (SIO) middleware development and support, interconnect technology development and deployment, SAN development and deployment, and archive.

The file systems element of the project provides end-to-end, high-performance networking and SIO infrastructure for the ASC program. Successfully meeting the ASC programmatic milestones requires carefully balanced environments in which the I/O infrastructure scales proportionally with increased ASC platform capabilities and application data needs. As the program moves toward exascale areas, these efforts will improve the scaling or programmability of the I/O in ASC applications for current and future large-scale machines. Current areas of investigation are Parallel Logged File System (PLFS), scalable indexing, burst buffer architectures, and scalable metadata.

Application Readiness capabilities are consolidated in this project, addressing issues with an application's production-run readiness on current and incoming computing systems at LANL. Working with subsystem teams such as systems management; file systems; and I/O, archive, and tools, the Application Readiness team identifies causes of unexpected behavior and deploys fixes in production so that system users are able to make productive use of the systems with their applications to solve their problems. The team provides production problem solving (create small problem reproducers, identify cause, consult with the relevant technical experts to find a solution, and verify the deployed solution), periodic stress testing/regression of production machines, new software version regression testing, system configuration verification and software stack deployment with real user applications and metrics, and analysis/profiling.

The project also includes software support capabilities focused on communication and networking libraries (MPI). The goal is to establish a strong development and analysis tool capability for current and next-generation HPC platforms, including parallel capabilities. The project is focused on working with the HPC tool community and vendors to identify, plan, and integrate tools into production environments and establish a solid support structure.

Required Capabilities

R1: Provide design and development of archival services, high-performance file systems, and I/O infrastructure for ASC program

R2: Provide high-performance networking and SIO infrastructure for the ASC program

R3: Improve the scaling or programmability of I/O in ASC applications for current and future systems (PLFS, scalable indexing, burst buffer, archive, and scalable metadata)

R4: Ensure application readiness on all LANL and ACES systems, with workloads performing at high efficiency and reliability

R5: Strong MPI development and support, with strategic/support plans and collaborations to continue to increase scalability and support of programming environment and analysis tools

Five-Year Plan

FY15

- Provide on-going support and testing for production file-systems and HPSS
- Deploy initially an open source or commercial archival solution in LANL production environment
- Provide application readiness support on commodity platforms, including MPI and threading models such as OpenMP
- Prototype and test a burst-buffer enablement library to insulate application developers from the complexities of next-generation file-systems, and schedulers

FY16

- Deploy I/O abstraction software for general use with the burst buffer
- Deploy a large-scale implementation of a future archive platform
- Collect public usage statistics for archive, burst buffer, and scratch file system
- Support users of LANL and ACES systems by tackling hard-to-diagnose problems, typically involving the interaction of applications with multiple aspects of the computational environment
- Provide MPI development and user support and sustain capabilities for scalable programming environment tools for current and future systems

FY17

- Refine I/O software based on usage statistics and lessons learned
- Publish collected statistics and participate in DOE
- Evaluate parallel file system performance for exascale machines
- Support users of LANL and ACES systems by tackling hard-to-diagnose problems, typically involving the interaction of applications with multiple aspects of the computational environment
- Provide communication library development and user support and sustain capabilities for scalable programming environment tools for current and future systems

FY18

- Plan and design parallel file system and archival storage strategy for exascale systems

- Support users of LANL and ACES systems by tackling hard-to-diagnose problems, typically involving the interaction of applications with multiple aspects of the computational environment
- Provide communication library development and user support and sustain capabilities for scalable programming environment tools for current and future systems

FY19

- Plan and design parallel file system and archival storage strategy for exascale systems
- Support users of LANL and ACES systems diagnosing problems, typically involving the interaction of applications with multiple aspects of the computational environment
- Provide communication library development and user support and sustain capabilities for scalable programming environment tools for current and future systems

Production Input/Output Services (SNL)

The Production I/O Services project represents SNL's participation in the DOE HPSS Consortium development project. HPSS provides the archival storage solution for ASC systems and is in direct alignment with ACES.

SNL's role in the HPSS project is to collaborate with tri-lab developers to design, implement, and test solutions that meet ASC requirements for all three labs.

Required Capabilities

R1: Current release levels of HPSS are stable and support existing capacity computing resources at Sandia

R2: HPSS performance enhancements to demonstrate sufficient file create rates and data ingestion speeds which will be needed to support future ATS; these are estimated to be on the order of 30,000 file creates per second, and ingest rates of a few TB/s.

R3: Data integrity to ensure no loss of data

Five-Year Plan

FY15

- Finalize HPSS version 7.5 deployment schedule to implement distributed core servers
- Design and prototype v8.1 with partitioned database across multiple systems, with a target of 30,000 file creates per second
- Proceed with development and testing of v8.1

FY16

- Deploy v8.1
- Design and prototype v8.2 with core services distributed across multiple systems, with a target of satisfying all capabilities with ~100,000 file creates per second

FY17

- Proceed with development and testing of v8.2

FY18

- Deploy v8.2

FY19

- Continue analysis and evaluation of HPSS performance and propose additional features to be developed as needed by the rapidly changing environment of next-generation computing systems and networks

Scalable Data Management (SNL)

The Scalable Data Management project provides critical R&D to support efficient and effective use of I/O capabilities on future exascale platforms. The research performed in this project addresses I/O concerns for application workflows, including efficient movement, management, and processing of data throughout the application workflow. This software will have tremendous impact on I/O for petascale and future systems because it allows for the creation of fully integrated scientific workflows that generate, analyze, and visualize data with minimal requirements for persistent storage. Current use of this technology includes data staging/caching to manage bursty I/O operations (for example, for checkpoints) and in-transit fragment detection for CTH, a co-design activity with the Scalable Data Analysis project.

Required Capabilities

R1: System software and libraries to support efficient data management and coupling of simulation and analysis

R2: Standard interfaces and programming models to leverage advanced architectures for data-intensive applications

Five-Year Plan

FY15

- Develop an I/O mini-app capable of representing SIERRA I/O patterns and demonstrating new capabilities developed by the scalable data management team
- Explore use of alternative storage architectures for SIERRA application workflows, in support of FY16 milestone to demonstrate analysis, visualization, and I/O capabilities for SIERRA

FY16

- Integrate and demonstrate in-memory storage capabilities (Kelpie) into SIERRA
- Investigate data-movement challenges for integrating advanced architectures into the application workflow

FY17

- Deploy data services with system support for placement in production environment

FY18

- Adapt data services to fully support code coupling, integrated analysis, and resilience in a prototype exascale environment

FY19

- Deploy frameworks for fully integrated application workflows on ATS

Scalable File Systems (SNL)

The Scalable File Systems project provides R&D to investigate and develop file systems for future exascale platforms. Issues of particular importance are extreme scalability and resilience. To address these issues, this project is developing a file system that decentralizes management of devices to support a high degree of heterogeneity within a system of inherently unreliable networks and storage devices. The central components of this peer-to-peer-like system are "smart" servers that have access to a variety of different local and remote media (for example, disk, non-volatile random access memory (NVRAM), memory, and tape) and are pervasive throughout the computing platform. These servers directly handle I/O requests, initiate third party transfers, or replicate the data as needed.

Required Capabilities

R1: Scalable, fault-tolerant, parallel file system

R3: Standard interfaces and programming models to leverage advanced architectures for data-intensive applications

Five-Year Plan

FY15

- Implement extended feature support for Sirocco using the peer-to-peer protocols
- Complete performance and correctness testing in preparation for deployment on Trinity

FY16

- Test and harden resilience capabilities of Sirocco file system

FY17

- Deploy data services with system support for placement in production environment

FY18

- Adapt file system to support next-generation ATS

FY19

- Deploy Sirocco for production use on broader ATS

Post-Processing Environments (WBS 1.5.4.6)

This level 4 product provides integrated post-processing environments to support end-user visualization, data analysis, and data management. The scope of this product includes planning, research, development, integration and deployment, continuing customer/product support, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include tools for metadata and scientific data management, as well as general-purpose and application-specific visualization, analysis, and comparison. Research includes innovative data access methods and visualization of massive, complex data—the use of open-source foundations will continue to be an important strategy for development of shareable advanced techniques. The product must develop solutions to address interactivity, scaling, tri-lab access for petascale platforms, and data analysis techniques needed to support effective verification and validation (V&V) and comparative analysis. Solutions for emerging platform architectures may in turn require customization and/or re-architecting of software to leverage hardware features. A continuing emphasis will be placed on tools for improving end-user productivity. The product also provides and supports infrastructure including office and collaborative space visualization displays, mechanisms for image data delivery, and graphics rendering hardware.

Scientific Visualization (LLNL)

The Scientific Visualization project conducts research and develops and supports tools for managing, visualizing, analyzing, and presenting scientific data. Research topics include topological analysis, particle visualization, and data compression techniques. Operational support for data analysis covers support of post-processing resources, including visualization servers, displays, and facilities. The visualization hardware architecture team engages in planning, test bed prototyping, testing of systems and components, and procurement and integration of new systems. Display efforts include support of high-resolution, high-performance display devices for theaters and collaborative use areas. The project installs, maintains, and consults on software visualization tools, and supports demonstrations on the PowerWalls. The project maintains unclassified and classified video production labs and consults on software such as resource management tools, movie players, animation, and visualization packages. The project exploits the latest capabilities of clustering hardware, GPUs, and parallel storage systems. Hardware capabilities include three production visualization servers and several PowerWall clusters. A video display infrastructure drives PowerWalls and smaller displays. Visualization researchers continued to perform work in areas of topology, compression, and advanced data analysis techniques.

Required Capabilities

R1: Data analysis hardware and software to support the analysis and post-processing of ASC simulation data

R2: PowerWalls for presenting high-resolution animations and images of large-scale data and operational support for high-level demonstrations

Five-Year Plan

FY15

- Evaluate, select, and deploy visualization hardware to support ASC data analysis needs
- Continue to maintain the data analysis and visualization hardware platforms and software environment and provide operational support for all visualization facilities, including supporting projection equipment and facilitating the use of the data analysis clusters and associated storage
- Support large-scale data analysis and visualization activities, including supporting ASC scientists with creation of visuals and movies for presenting and analyzing scientific data
- Exploit research results in data analysis and visualization for ASC simulations, including data compression, topological methods, and optimization of streamline tracing algorithm

FY16

- Maintain the data analysis and visualization hardware platforms and software environment and provide operational support for visualization facilities
- Refresh PowerWall driver technology as needed

FY17

- Maintain the data analysis and visualization hardware platforms and software environment and provide operational support for visualization facilities
- Support large-scale data analysis and visualization activities

FY18

- Maintain the data analysis and visualization hardware platforms and software environment and provide operational support for visualization facilities
- Support large-scale data analysis and visualization activities
- Gather requirements, procure, and support data analysis and visualization platform to support the 2017 Sierra ATS

FY19

- Maintain the data analysis and visualization hardware platforms and software environment and provide operational support for visualization facilities
- Refresh PowerWall technology as needed
- Support large-scale data analysis and visualization activities, including support for GPUs

Scientific Workflow and Data Management (LLNL)

The Scientific Workflow and Data Management project provides users with powerful and time-conserving ways to access, search, compare, and archive large-scale scientific data, and new high-level tools for managing the simulation workflow. This is achieved through the development of production-quality applications that enhance data management capabilities and the creation of innovative interfaces to job monitoring and vertical application frameworks.

Hopper and Chopper are the principal products of the data management effort. In the simulation workflow area, the Lorenz Web-based HPC application suite forms a foundation for providing new ASC-specific capabilities. Lorenz uses advanced Web technologies to make HPC more accessible, saving the user time while also helping the resources to be used more effectively.

Required Capabilities

R1: Flexible data management software that helps users make efficient and appropriate use of resources

R2: Highly customizable and easily accessible user interface to HPC resources and information

Five-Year Plan

FY15

- Release new versions of Hopper and Chopper with a focus on incorporating MPI and non-MPI based parallelism for copy and directory operations
- Incorporate RobinHood-based Lustre metadata in Hopper and Chopper, significantly improving disk usage and related scanning operations for users
- Extend the MyLC dashboard to include more complete information about changes occurring within the center, and provide users with a variety of ways to subscribe to this information
- Investigate workflow-related frameworks and techniques and prototype elements of an identity management replacement tool

FY16

- Explore new protocols and techniques for managing extreme scale data management

FY17

- Incorporate support for the dominant exascale data storage and handling protocols into high-level tools for end users

FY18

- Extend Lorenz to support advanced features of next-generation resource managers

FY19

- Implement and coordinate workflows that are tuned to exascale, allowing users to leverage the work of domain experts

Visualization and Data Analysis (LANL)

Data analysis and visualization are key capabilities in taming and understanding the ever-increasingly large datasets generated from extreme-scale scientific simulations. This project comprises research, development, deployment of software and facilities to production and ongoing expert support in this.

The production and facilities component of the project is to provide LANL weapons designers with visualization systems research and support, and to provide analytic expertise to help LANL weapons designers utilize the full power of the hardware and software infrastructure for visualization and data analysis developed and deployed by ASC, thus improving the physics understanding of their weapons simulations.

Technical staff members funded by the project assist in the design theater and the co-laboratories. The project also deploys within the design community in X Division a small group of individuals with expert knowledge in both visualization and weapons science to work directly with the designers. Capabilities include the design and deployment of new visualization systems, briefing support, and support of large facilities, such as the CAVE and the PowerWall. The project provides assistance, training, and developing of new tools to work with these facilities. Development, deployment, and maintenance of any needed visualization corridor software are also provided by this project.

The project is responsible for the EnSight visualization and data analysis software, including maintaining the EnSight software installation laboratory-wide, providing local user support in the use of the software, and acting as a bridge between the LANL design community and the EnSight developers at Computational Engineering International.

The R&D aspect of the project develops new visualization algorithms and systems to meet current and future capability requirements for ASC simulations. This work is required to address ASC workloads: massive data sizes, complex results, and the use of unique supercomputing architectures.

ASC simulations are currently producing massive amounts of data that threaten to outstrip the ability to visualize and analyze it. Therefore, it is important to understand how to triage data within the simulation as it is generated using techniques such as in-situ analysis for data reduction and visualization kernels that run on the supercomputing platform, including data analysis, visualization, and rendering methods. Continue to integrate/demonstrate situ techniques in IC, the project will extend this capability to other codes and to enabling the design community to explore the new level of analytic flexibility now available.

Another important analysis mode involves storing reduced data for later interactive analysis and visualization (that is, post-processing analysis). Extreme-scale ASC databases, petabytes or larger, already exist and are growing in number and size. LANL's long-term objective is to build a post-processing analysis system that, within reasonable limits, can manage performance optimization tasks automatically, be fault-tolerant, provide a set of high-level serial and parallel programming primitives for carrying out complex queries and computations on the stored data, and provide fast enough execution

to enable interactive, iterative discovery. Current approaches range from the database-driven (structured query language (SQL)) to parallel computation-driven (Hadoop/MapReduce), neither of which alone can meet the needs of the target community. LANL will build on these technologies, integrating the best of both approaches.

Required Capabilities

R1: Visualization computing systems and facilities operating at high availability

R2: Support for visualization computing in the advanced and production computing systems, and expert visualization support for ASC simulation scientists

R3: *In-situ* integration with scientific simulation codes

R4: Quantifiable data reduction and triage

R5: Portable, scalable visualization and analysis software and hardware infrastructure

Five-Year Plan

FY15

- Continue to support and maintain production visualization systems and production visualization tools, including directly with designers, and to support the production integration of visualization hardware for Trinity and CTS
- Move R&D tech transfer to production in-situ, improving feature set and statistics for in situ data analysis framework to additional ASC codes
- Continue development of PISTON portable hardware-accelerated visualization library
- Continue expert technical analysis of visualization and big data for ASC

FY16

- Continue to support and maintain production visualization systems and to promote new discoveries in weapons science by advanced applications of visualization and data analysis
- Extend PISTON to include a sufficiently comprehensive suite of commonly used visualization and analysis operators and a sufficiently accessible user interface such that it can be used by domain experts in a wide array of scientific applications
- Explore, develop, and deploy quantitative visualization and analysis for ensembles, coupled multiscale, and multivariate physics

FY17

- Continue to support and maintain production visualization systems and to promote new discoveries in weapons science by advanced applications of visualization and data analysis

- Deliver an integrated software infrastructure for visualization and analysis applied to ASC codes of interest as part of a Level 2 milestone

FY18

- Continue to support and maintain production visualization systems and to promote new discoveries in weapons science by advanced applications of visualization and data analysis
- Explore data analysis and visualization technologies for and integrated hardware and software infrastructure for extreme-scale systems

FY19

- Continue to support and maintain production visualization systems and to promote new discoveries in weapons science by advanced applications of visualization and data analysis
- Explore data analysis and visualization technologies for and integrated hardware and software infrastructure for extreme-scale systems via in situ and big data paradigms

Scalable Data Analysis (SNL)

The Scalable Data Analysis project provides data analysis capabilities and support for a range of SNL ASC customers—from analysts and code developers to algorithm designers and hardware architects. Capabilities include data manipulation, data transformation, and data visualization that contribute to insight from computational simulation results, experimental data, and/or other applicable data. A project emphasis is to deliver and support scalable capabilities that support increasing data sizes, data sources, and platform processor counts for ASC complex applications and system architecture.

This project includes production deployment and support services that enable ASC customers to carry out data analysis on ASC systems. This includes porting and installation of tools onto production systems; maintenance, testing, debugging, refinement and integration of tools in the end-to-end system environment as needed to assure effective end-user capabilities; and user support. SNL priorities include a focus on delivering and supporting analysis capability for Cielo and subsequent ACES platforms.

Current tools include scalable data analysis software released open source through ParaView and the Visualization Toolkit (VTK), an early-release in-situ data analysis library (Catalyst) for coupling directly with running simulation codes, and R&D prototypes for the analysis of results from ensembles of simulation runs. Current hardware platforms for data analysis are limited to data analysis/visualization partitions on the compute platforms with an emphasis on delivery of visualizations to desktop.

Partnering with ASC customers and other product areas, this project will continue to build on its successful ParaView and VTK-based products. The project performs R&D that advances these capabilities as needed for evolving next-generation architectures, ensuring that ASC's investment in data analysis and visualization will provide advanced capabilities on platforms from Cielo through future exascale systems.

Required Capabilities

R1: Scalable post-processing tools for analysis and visualization on existing and next-generation architectures

R2: Scalable in-situ data analysis and visualization capability to mitigate impact of I/O mismatch and power constraints at extreme scales

R3: Dedicated hardware resources for interactive data analysis and visualization

R4: Production deployment and support for interactive applications, in-situ analysis capabilities, and infrastructure

Five-Year Plan

FY15

- "Harden" the Catalyst in situ capability through deep evaluation of overheads, including memory and performance

- Begin integration of analysis and visualization capabilities with the Sandia Analysis Workbench (in support of FY16 and FY17 Level 2 milestones related to application workflow demonstrations)
- Prototype integration of multi/many core code into production tools
- Continue *ParaView* and Catalyst releases, with production support
- Deliver scalable analysis and visualization capabilities for Cielo and Sequoia customers

FY16

- Deliver analysis and visualization tools on Trinity for Limited Availability
- Perform experiments of multi-mode in-situ processing at scale on Trinity
- Deliver scalable analysis and visualization capabilities for Cielo and Sequoia customers

FY17

- Continue ParaView and Catalyst releases, fully transitioned to multi/many core
- Release pre-production versions of post processing tools for ensemble and in-situ analysis
- Demonstrate and evaluate use of analysis and visualization capabilities for SIERRA applications at large scale on ATS in support of a Level 2 milestone
- Evaluate custom adaptations of analysis capabilities for other applications (for example, Albany, Alegra, and XYCE)

FY18

- Release production versions of post-processing tools for ensemble and in-situ analysis

FY19

- Production release of workflow-enabled tools for in situ and post-processing analysis and visualization

Facility Operations and User Support (WBS 1.5.5)

This sub-program provides both necessary physical facility and operational support for reliable, cross-lab production computing and storage environments as well as a suite of user services for effective use of ASC tri-lab computing resources. The scope of the facility operations includes planning, integration and deployment, continuing product support, software license and maintenance fees, procurement of operational equipment and media, quality and reliability activities, and collaborations. FOUS also covers physical space, power and other utility infrastructure, and local area network (LAN)/wide area network (WAN) networking for local and remote access, as well as requisite system administration, cyber-security, and operations services for ongoing support and addressing system problems. Industrial and academic collaborations are an important part of this sub-program.

Collaborations (WBS 1.5.5.3)

This level 4 product provides programmatic support for collaboration with external agencies on specific HPC projects. This product also includes collaborations with internal or external groups that enable the program to improve its planning and execution of its mission.

Program Support (LLNL)

The Program Support project provides service to the ASC program. Program Support services include procurement and contracting, project management, and meeting support. These services are in support of both tri-lab and LLNL-only activities, including collaborations with academic, industrial, and other government agencies.

Required Capabilities

R1: Support procurement, planning, collaboration, meetings, and contracting needs of NNSA tri-lab ASC program

Five-Year Plan

FY15

- Continue FY15 procurement, contract management, and program planning needs
- Support annual HPC Operations Review meeting with Office of Science
- Support bi-annual PSP meetings
- Provide contract management and procurement support of tri-lab procurements and system deliveries

FY16

- Continue FY15 procurement, contract management, and program planning needs
- Support annual HPC Operations Review meeting with SC
- Support bi-annual PSP meetings
- Manage delivery of CTS systems to tri-labs
- Manage Critical Decision (CD) process for 2017 ATS

FY17

- Continue FY16 procurement, contract management, and program planning needs
- Continue contract management for delivery of CTS to tri-labs
- Manage CD process and HQ reporting for the Sierra ATS, including system delivery

FY18

- Continue FY17 procurement, contract management, and program planning needs
- Support annual HPC Operations Review meeting with SC
- Support bi-annual PSP meetings
- Write the Sierra ATS CD4 and final Sierra ATS HQ reports

FY19

- Continue FY18 procurement, contract management, and program planning needs

- Support annual HPC Operations Review meeting with SC
- Support bi-annual PSP meetings

Program Support (LANL)

Through the Program Support project, LANL provides support to the national program, both by providing resources and expertise to the Federal program office and by participating in coordination and integration activities for the tri-lab program.

Required Capabilities

R1: Support ASC Principal Investigator (PI) meeting

R2: Host PSP meeting

R3: Publish ASC eNews online newsletter

R4: Support PSAPP

R5: Tri-lab publications

Five-Year Plan

FY15

- Support PSAAP-II collaborations
- Host PSP meeting
- Publish ASC eNews online newsletter
- Participate in PSP and PI meeting
- Provide LANL support for HQ

FY16

- Participate in PI and PSP meetings
- Publish ASC eNews online newsletter
- Provide LANL support for HQ

FY17

- Host PSP Meeting
- Publish ASC eNews online newsletter
- Provide LANL support for HQ

FY18

- Participate in PI and PSP meetings
- Publish ASC eNews online newsletter
- Provide LANL support for HQ

FY19

- Host PSP meeting

- Publish ASC eNews online newsletter
- Provide LANL support for HQ

Program Support (SNL)

The Program Support project provides critical coordination and integration activities essential to the success of ASC. It is divided into two distinct parts: 1) provide ASC programmatic planning, reviews, and communications; and 2) provide SNL outreach to the other institutions and programs.

This capability is critical to the ASC SNL program integration, communication, and management within the laboratory and with the external community. A significant management and integration function in this project is captured in the Science Applications International Corporation (SAIC) contract that provides support for NNSA HQ and SNL in communications and logistics. External advisory boards supported through this project also provide feedback to the ASC leadership team regarding the maturation of the predictive engineering sciences capability and the quality of SNL's computational science R&D. Support of external collaborations, including PSAAP and the exascale initiative (with DOE/SC), is also included in this project.

Required Capabilities

R1: Management of ASC program cost, schedule, and performance

R2: Management of SAIC to provide administration support for ASC HQ

R3: External reviews of program-supported predictive capability development

R4: Engagement with ASC federal managers and tri-lab ASC executives

R5: Technical engagement and program planning for the PSAAP

R6: Technical engagement and program planning for the proposed DOE Exascale Initiative

Five-Year Plan

FY15

- Organize and host sixth PESP meeting
- Support external review panel meetings for Qualification Alternatives to the Sandia Pulsed Reactor (QASPR), the Engineering Sciences External Advisory Board, and the Computer and Information Sciences External Advisory Board
- Support programmatic needs of the PSAAP II program
- Complete lead lab responsibilities for SC14 Conference
- Support programmatic needs of NNSA tri-lab ASC program and ASC executive committee

FY16

- Organize and host seventh Predictive Engineering Science Panel meeting

- Support external review panel meetings for QASPR, the Engineering Sciences External Advisory Board, and the Computational Sciences External Advisory Board
- Support programmatic needs of the PSAAP II program
- Manage the SAIC contract to provide various administration support for HQ
- Support programmatic needs of NNSA tri-lab ASC program and ASC executive committee

FY17

- Organize and host eighth Predictive Engineering Science Panel meeting
- Organize and host ASC Principal Investigator's Meeting
- Organize attendance, booth, and meeting logistics for SC17
- Support external review panel meetings for QASPR, the Engineering Sciences External Advisory Board, and the Computational Sciences External Advisory Board
- Support programmatic needs of the PSAAP II program
- Manage the SAIC contract to provide various administration support for HQ
- Support programmatic needs of NNSA tri-lab ASC program and ASC executive committee

FY18

- Organize and host ninth Predictive Engineering Science Panel meeting
- Complete lead lab responsibilities for SC17 conference
- Support the Engineering Sciences External Advisory Board, and the Computational Sciences External Advisory Board
- Support programmatic needs of the PSAAP II program
- Manage the SAIC contract to provide various administration support for HQ
- Support programmatic needs of NNSA tri-lab ASC program and ASC executive committee

FY19

- Organize and host the tenth Predictive Engineering Science Panel meeting
- Support the Engineering Sciences External Advisory Board and the Computational Sciences External Advisory Board
- Support programmatic needs of the PSAAP II program (potential)
- Manage the SAIC contract to provide various administration support for HQ
- Support programmatic needs of NNSA tri-lab ASC program and ASC executive committee

Applications in Support of Manufacturing Production and Connectivity (Y-12)

The Applications in Support of Manufacturing Production and Connectivity project supports the utilization of ASC codes and computing resources to solve production manufacturing problems through modeling and simulation. The project includes support for connecting to ASC computing resources and job submission, execution, and visualization. The project provides the infrastructure necessary to test applications and scenarios before deployment on larger ASC resources. Development and deployment of software to support the solution of manufacturing problems is also supported by the project. Visualization techniques that can be utilized in the Y-12 network and computing infrastructure will be evaluated and implemented. Finally, participation in Nuclear Weapons Complex (NWC) ASC-related activities is covered.

Required Capabilities

R1: Local cluster computer capable of code development and performing small to mid-scale analysis of radiation transport, heat transfer, and image, and signal analysis problems

R2: Access to commercial and ASC-developed software for heat transfer, fluid flow, and solid mechanic (finite-element) analysis

R3: Support for development of models to solve production manufacturing problems and visualization techniques that can be utilized in manufacturing environment

Five-Year Plan

FY15

- Update the ART optical tracking system with a 3D fly stick to manipulate images; develop a virtual training demonstration with immersive 3D stereo and head and fly-stick tracking
- Demonstrate collaborative viewing of a stereo immersive scene in Visionary Render
- Upgrade the immersive environment with high-resolution stereo projection
- Develop specifications to upgrade the classified simulation modeling terminal servers
- Develop a dynamic virtual factory simulation of the proposed Lithium Production Facility

FY16

- Implement the capability for team participation in virtual reality training scenarios
- Upgrade simulation modeling computing infrastructure to support manufacturing simulation and production modeling
- Continue to collaborate with other sites (NWC colleagues) to best utilize codes and deploy capabilities as necessary to enable collaboration

- Continue to evaluate new codes on the Y-12 cluster and utilize Y-12 and remote ASC cluster resources to solve production manufacturing problems
- Participate in National Security Enterprise ASC activities

FY17

- Continue to collaborate with other sites (NWC colleagues) to best utilize codes and deploy capabilities as necessary to enable collaboration
- Continue to evaluate new codes on the Y-12 cluster and utilize Y-12 and remote ASC cluster resources to solve production manufacturing problems
- Participate in National Security Enterprise ASC activities
- Upgrade simulation modeling computing infrastructure to support manufacturing simulation and production modeling

FY18

- Continue to collaborate with other sites (NWC colleagues) to best utilize codes and deploy capabilities as necessary to enable collaboration
- Continue to evaluate new codes on the Y-12 cluster and utilize Y-12 and remote ASC cluster resources to solve production manufacturing problems
- Participate in National Security Enterprise ASC activities
- Upgrade simulation modeling computing infrastructure to support manufacturing simulation and production modeling

FY19

- Continue to collaborate with other sites (NWC colleagues) to best utilize codes and deploy capabilities as necessary to enable collaboration
- Continue to evaluate new codes on the Y-12 cluster and utilize Y-12 and remote ASC cluster resources to solve production manufacturing problems
- Participate in National Security Enterprise ASC activities
- Upgrade simulation modeling computing infrastructure to support manufacturing simulation and production modeling

System and Environment Administration and Operations (WBS 1.5.5.4)

This product provides requirements planning, initial deployment, configuration management, and ongoing operational support for reliable production computing and storage environments. Activities include system and network administration and operations, user support, hardware maintenance, licenses, and common tri-lab computing environment integration and support.

System and Environment Administration and Operations (LLNL)

This project provides necessary operational support for reliable production computing environments. The following activities are included: system administration and operations, software and hardware maintenance, licenses and contracts, computing environment security and infrastructure, requirements planning, initial deployment, production computing services, and tri-lab system integration and support. Included within the scope of this product is the operational support for systems used as part of partnerships with academic, industrial, and other governmental agencies.

Required Capabilities

R1: System configuration

R2: System and user security development and support

R3: System administration and monitoring

R4: File system support

R5: 24x7 system monitoring and user support

R6: Self hardware maintenance of all platforms

R7: Logistics coordination and spares management

Five-Year Plan

FY15

- Continue to deploy Virtual Machine based infrastructure servers and migrate services to new platforms
- Retire Juno and Graph
- Migrate NFS home directories to new hardware
- Integrate identity lifecycle policies and capabilities with institutional data sources
- Investigate options for replacing the identity management system (IDM) workflow, approvals, and provisioning processes
- Establish distinct development, pre-production, and deployment environments for the security infrastructure
- Implement log-based security event analysis and detection

FY16

- Deploy follow-on CTS-1 systems
- Perform ongoing system administration and operation of production platforms and file systems
- Provide 24x7 operations and monitoring of HPC computing resources

- Provide hardware self-maintenance for current and future ASC platforms
- Deploy 2016 early delivery ATS
- Perform ongoing system and user security development and support

FY17

- Perform ongoing system administration and operation of production platforms and file systems
- Provide 24x7 operations and monitoring of HPC computing resources
- Provide hardware self-maintenance for current and future ASC platforms
- Perform ongoing system and user security development and support

FY18

- Deploy 2017 ATS
- Perform ongoing system administration and operation of production platforms and file systems
- Provide 24x7 operations and monitoring of HPC computing resources
- Provide hardware self-maintenance for current and future ASC platforms
- Perform ongoing system and user security development and support

FY19

- Perform ongoing system administration of production platforms and file systems
- Perform ongoing system and user security development and support
- Provide 24x7 operations and monitoring of HPC computing resources
- Provide hardware self-maintenance for current and future ASC platforms

Hotlines and System Support (LLNL)

The Hotlines and System Support project provides users with a suite of services enabling effective use of ASC computing resources for the tri-lab as well as academic and industrial collaborations. This project includes computer center hotline and help desk services, account management, Web-based system documentation, system status information tools, user training, incident management systems, and application analyst support. Services are provided to both LLNL users as well as users from external sites, including LANL, SNL, and the ASC Alliance sites.

This project provides accounts administration, technical consulting, and documentation and training to facilitate the effective use of LLNL HPC systems. An accounts specialist team provides all account management services necessary for users to obtain accounts and access LLNL HPC systems. This includes account creation and removal, bank allocations, token management and visitor tracking for foreign national users. The technical consultant team provides technical support to LLNL users to enable their effective use of LLNL HPC systems. Consulting services vary from helping new users configure their environment, assisting experienced users with optimization of codes, and supporting other Livermore Computing (LC) staff with monitoring of file systems, batch queues, and user environments. Extensive Web documentation, user manuals, technical bulletins, and training are provided to users via email, Web, and in-person training.

Required Capabilities

R1: Provide user account management, bank allocations, and One-Time Password (OTP) support to all users of LLNL supercomputers

R2: Provide technical consulting services to all users of LLNL supercomputers

R3: Provide current technical documentation and training materials, including courses on the use of LLNL supercomputers

R4: Create and maintain tools for LLNL users that facilitate the effective use LLNL supercomputers

Five-Year Plan

FY15

- Continue to provide ongoing support services for hotline operations, documentation, and training
- Continue to migrate, reorganize, and consolidate existing user documentation to the LC Confluence server
- Learn about accelerator technologies and their software libraries (GPUs, MIC) in preparation for providing support to the user community in their use
- Provide trusted agent support to for National Security Systems Public Key Infrastructure (PKI) on the Enterprise Secure Network

- Become familiar with the Service Now Incident Management software in preparation for migration to it from Front Range in FY16

FY16

- Assist users in the migration of applications to the new CTS-1 systems
- Develop training and documentation for the early delivery ATS-2 system
- Deploy the new knowledge base of information for LC staff and users

FY17

- Develop training and documentation for the ATS-2 system
- Assist users in the migration of applications to the early delivery ATS-2 system

FY18

- Continue to provide ongoing support services for hotline operations, documentation, and training
- Assist users in the migration of applications to the ATS-2 system

FY19

- Continue to provide ongoing support services for hotline operations, documentation, and training
- Develop training and documentation for the CTS-2 systems

Facilities, Network, and Power (LLNL)

The Facilities, Network, and Power project provides for the necessary physical facilities, utilities, and power capabilities to ASC systems. Work in this area includes adequate raised floor space, flexible cooling solutions, and power to site large-scale ASC platforms. In addition, this project funds needed office, meeting room, and auxiliary space to enable a highly motivated and effective staff. Also included are classified and unclassified facility networks, wide-area classified networks, and ongoing network operations. This project also enables enhanced collaborations with academic and industrial partners.

Required Capabilities

R1: Provide support for electrical, mechanical, cooling, and network services infrastructures

R2: Manage ASC computing facility upgrades

R3: Support network backbones, hardware and software, in LLNL collaboration zone (CZ), restricted zone (RZ), and secure computing facility (SCF) networks

R4: Support DisCom WAN configurations and ESNet

Five-Year Plan

FY15

- Begin construction of the new unclassified HPC facility to house unclassified systems, including the next-generation CTS-1 clusters
- Begin to purchase long lead electrical equipment to house electrical equipment that allows for more widely varying voltages and block load in preparation for the 2017 system
- Complete commissioning of the B453 chilled-water system cutover project
- Relocate all computational equipment from B115 to other enduring HPC facilities, eliminating ongoing building maintenance costs to program; continue to enhance diagnostic and monitoring of IB fabrics on IB-attached Lustre file systems
- Continue to analyze and evaluate emerging network technologies

FY16

- Provide ongoing operations and maintenance of electrical and mechanical systems for ASC computing facilities
- Begin ATS facilities preparation project
- Provide design, technical support, configuration management, technical deployment expertise for all system and components of the LC's networking infrastructure

FY17

- Commission facilities for 2017 ATS
- Provide ongoing operations and maintenance of electrical and mechanical systems for ASC computing facilities
- Provide design, technical support, configuration management, technical deployment expertise for all system and components of the LC's networking infrastructure

FY18

- Provide ongoing operations and maintenance of electrical and mechanical systems for ASC computing facilities
- Provide design, technical support, configuration management, and technical deployment expertise for all system and components of the LC's networking infrastructure

FY19

- Begin CTS-2 facilities preparation project
- Provide ongoing operations and maintenance of electrical and mechanical systems for ASC computing facilities
- Provide design, technical support, configuration management, and technical deployment expertise for all system and components of LC's networking infrastructure

System Administration and Storage (LANL)

The System Administration and Storage project covers all services for computational systems operated by LANL for the purpose of providing an HPC production computing environment for weapons designers, developers, and engineers. The project works with users to troubleshoot problems experienced while running their applications, and helps users transition from old to new computing platforms. The capabilities include system configuration, system and user security, resource management, system administration and monitoring, archival storage, Panasas, and NFS.

Required Capabilities

R1: System configuration

R2: System and user security

R3: Resource management

R4: System administration and monitoring

R5: Archival storage

R6: File system support

Five-Year Plan

FY15

- Support HPC systems by conducting ongoing daily system and storage administration with continuous monitoring of production systems and infrastructure servers
- Provide support for Trinity (ATS-1) and CTS-1 site preparation project
- Provide support for CTS-1 integration and production readiness

FY16

- Support HPC systems by conducting ongoing daily system and storage administration with continuous monitoring of production systems and infrastructure servers
- Retire Luna, Moonlight, and Cielo
- Provide support for ASC/HPC systems/architectures
- Provide support for Trinity integration readiness milestone

FY17

- Support HPC systems by conducting ongoing daily system and storage administration with continuous monitoring of production systems and infrastructure servers
- Provide support for ASC/HPC systems/architectures
- Provide support for Trinity production readiness milestone

FY18

- Support HPC systems by conducting ongoing daily system and storage administration with continuous monitoring of production systems and infrastructure servers
- Provide support for ASC/HPC systems/architectures

FY19

- Support HPC systems by conducting ongoing daily system and storage administration with continuous monitoring of production systems and infrastructure servers
- Provide support for ASC/HPC systems/architectures

Operations and Procurement Support (LANL)

The Operations and Procurement Support project provides around-the-clock operations and monitoring of the scientific computing resources, including performance computers such as Luna, Typhoon, ViewMaster II, Cielo, Moonlight, and data storage and retrieval systems such as the HPSS. In addition to monitoring all components 24x7x365, the computer operators provide systems hardware maintenance for all ASC platforms. Working with the vendor system engineers, the operators also provide backup hardware support for the Cielo capability system. This includes all components of the production computing environment, from compute engines, hardware, fileservers, archival storage systems, the facilities they reside in and utilities they are dependent upon, to all required software on these systems.

The procurement support aspect of this project assists customers with the technical and administrative aspects of planning, procurement, and contractual agreements for computer hardware and software products and services.

Required Capabilities

R1: Provide 24x7x365 operations and monitoring of ASC/HPC resources

R2: Provide hardware maintenance for ASC/HPC resources

R3: Provide procurement support for ASC/HPC hardware, software, services, and contracts

Five-Year Plan

FY15

- Provide 24x7x365 operations and monitoring of ASC/HPC computing resources
- Develop reports for projecting out-year hardware maintenance costs for current and future ASC platforms
- Provide technical and administrative support for procurement of ASC/HPC platforms, supporting facility hardware, software, and other products and services required by ASC/HPC

FY16

- Provide 24x7x365 operations and monitoring of HPC computing resources
- Provide hardware self-maintenance for current and future ASC platforms
- Decommission Luna and Moonlight ASC Platforms
- Provide system administration, network administration, and Zenoss interface development on monitoring and hardware testing infrastructure
- Provide technical and administrative support for procurement of HPC platforms

FY17

- Provide 24x7x365 operations and monitoring of HPC computing resources
- Provide hardware self-maintenance for current and future ASC platforms
- Provide system administration, network administration, and Zenoss interface

FY18

- Provide 24x7x365 operations and monitoring of HPC computing resources
- Provide hardware self-maintenance for current and future ASC platforms
- Provide system administration, network administration, and Zenoss interface
- Provide technical and administrative support for procurement of HPC platforms

FY19

- Provide 24x7x365 operations and monitoring of HPC computing resources
- Provide hardware self-maintenance for current and future ASC platforms
- Provide system administration, network administration, and Zenoss interface
- Provide technical and administrative support for procurement of HPC platforms

Computing Platform Integration and Deployment (LANL)

The scope of the Computing Platform Integration and Deployment project is to accept delivery and begin deployment of production CTS platforms. This includes participating in developing the design requirements as part of a tri-lab requirements planning team. Primary capabilities include completing the acceptance tests, diagnostics test, integrating the systems into the LANL unclassified network, system stabilization, and transition into the classified network. Included in this project is support for the ASC CTS acquisition strategy and provision for requirements that help to achieve the strategy.

The objective of the project is the integration of all hardware and software components to deliver a system environment to application users for programmatic work. This includes site preparation to prepare the Simulation and Computing Complex (SCC) facility for deploying these production capacity systems. The integration and deployment activities will focus on the following areas: System/OS, File Systems, Interconnect, External Network including PaScalBB, Regression Testing, Monitoring, and Application Readiness.

Required Capabilities

R1: Platform acceptance and diagnostic testing

R2: Platform integration into LANL unclassified and classified networks

R3: System stabilization

Five-Year Plan

FY15

- Deploy CTS-1 at Los Alamos
- Develop plan and schedule for system accreditation

FY16

- Provide production support for CTS-1 system

FY17

- Provide production support for CTS-1 systems

FY18

- Develop requirements, with tri-lab design team, for CTS-2
- Provide production support for CTS-1 systems

FY19

- Deploy CTS-2 at LANL
- Develop plan and schedule for system accreditation
- Provide production support for CTS-2 systems

Integrated Computing Network Consulting, Training, Documentation, and External Computing Support (LANL)

The Integrated Computing Network Consulting, Training, Documentation, and External Computing Support project is responsible for direct customer service for local and remote users of ASC/LANL resources, the development and delivery of documentation and training materials for ASC/LANL resources, usage statistics, and an administrative interface for the ASC tri-lab, Alliance users, and other external ASC/HPC users. The primary capabilities consist of user support services, operational metrics for an HPC environment on, for example, usage and availability, Web-page development to present this information to system personnel and users, and the development of user documentation and training.

Required Capabilities

R1: Ongoing user support services

R2: Operational metrics for HPC resource usage and availability

R3: Development of user documentation and training

R4: Web-page development providing access to HPC services

Five-Year Plan

FY15

- Prepare Trinity (ATS-1) system documentation and training materials
- Develop CTS-1 system documentation and training materials
- Provide ongoing consulting and user support services, documentation, and training for ASC platforms and architectures

FY16

- Provide assistance for use of emerging ASC platforms/architectures

FY17

- Continue to enhance user support tools, capabilities, and infrastructure available to users and the user support team
- Provide assistance for use of emerging ASC platforms/architectures

FY18

- Develop next-generation system documentation and training materials
- Continue to enhance user support tools, capabilities, and infrastructure available to users and the user support team

FY19

- Develop CTS-2 system documentation and training materials

- Continue to enhance user support tools, capabilities, and infrastructure available to users and the user support team
- Develop ATS-3 system documentation and training materials

Facilities, Networking, and Power (LANL)

The Facilities, Networking, and Power project is responsible for the engineering, design, operation, and maintenance of the mission-important electrical, mechanical, cooling, network services, and other computing infrastructure in support of the ASC program. The project provides support for infrastructure design upgrades, project and space management, user interface and oversight, demolition and decommissioning of older systems, network backbones, user LANs, classified/unclassified network hardware and services, DisCom WAN, and computer site preparation for new platforms. Because the tri-lab community requires the systems to be operational at all times, the project provides on-call support after hours and on weekends for facility related issues.

Required Capabilities

R1: Provide support for electrical, mechanical, cooling, and network services infrastructures

R2: Manage ASC computing facility upgrades

R3: Provide decommissioning services for resources reaching end-of-life

R4: Support network backbones, hardware and software, in LANL yellow and red networks

R5: Support DisCom WAN

Five-Year Plan

FY15

- Provide support for Trinity (ATS-1) and CTS-1 platform integration into SCC computing facility
- Provide ongoing operations and maintenance of electrical and mechanical systems for ASC computing facilities
- Integration of Next Generation Backbone (NGBB) technology in support of Trinity (ATS-1) and CTS-1 platforms

FY16

- Provide ongoing operations and maintenance of electrical and mechanical systems for ASC computing facilities
- Design SCC upgrade for mechanical and electrical infrastructure project for ATS-3
- Complete NGBB build and integration

FY17

- Provide ongoing operations and maintenance of electrical and mechanical systems for ASC computing facilities

FY18

- Provide ongoing operations and maintenance of electrical and mechanical systems for ASC computing facilities

FY19

- Provide support for CTS-2 platform integration into SCC computing facility
- Provide ongoing operations and maintenance of electrical and mechanical systems for ASC computing facilities
- Provide ongoing operations and maintenance of electrical and mechanical systems for ASC computing facilities

Production Computing Services (SNL)

The Production Computing Services project's goals are to operate and maintain all ASC production platforms and associated support systems, and operate data services and visualization systems, long-term hierarchical storage services, high-performance network systems, tri-lab compatible cyber authentication and authorization systems, and monitoring and reporting services. This project supports tri-lab capability platform resource allocations and coordinates with tri-lab peers in establishing priority scheduling, if required. This project coordinates the integration and deployment of TLCC and CTS capacity systems into SNL's production computing environment, in collaboration with WBS 1.5.5.6 Common Computing Environment. Support of CCE common service and environment decisions and configuration management activities are also provided.

This project has expertise in operating capacity computing clusters; integrating file servers at the system or facility-wide level; deploying new computing, storage, and data management platforms; and in retiring end-of-life platforms. System administration for complex HPC environments is provided, as are design and development activities for new innovative advanced architecture computing platforms.

Required Capabilities

R1: Computing facility space, power, and cooling sufficient for current and future platforms

R2: Experienced technical staff and support personnel versed in computing systems (both hardware and software), networks and protocols (interconnects and local networks), file systems and batch management

R3: Wide Area High Speed network design and operations support

R4: Data transfer and storage systems and support personnel

R5: Tier 2 and Tier 3 support for trouble resolution of individual systems and combined services

Five-Year Plan

FY15

- Begin construction plans for new computing facility
- Begin procurement activity for next-generation capacity systems
- Perform ACES integration activities for Trinity; engage with the Lawrence Berkeley National Laboratory NERSC production operations team in implementing LDMS
- Continue operations of all production systems

FY16

- Complete installation and acceptance of Capability Class NSCC resource

- Integrate expansion of file systems and data archive platforms sized to service the Capability Class system
- Continue operations of all production systems

FY17

- Integrate next generation of CTS-1 capacity compute platforms into ASC and NSCC environments
- Continue operations of all production systems

FY18

- Continue operations of all production systems

FY19

- Continue operations of all production systems
- Plan for, acquire, and integrate file systems and archive systems to support current systems and anticipated growth of next CTS-2 platforms

User Support (SNL)

The User Support project provides user support and associated resources for SNL computing systems and tri-lab resources. User support activities focus on improving the productivity of the entire user community, local and remote, in utilizing the ASC HPC resources.

This project deploys and maintains the following SNL capabilities for user support: 1) coordination between user support activities and leadership in adopting Information Technology Infrastructure Library (ITIL) principles and practices; 2) ITIL-compliant incident, problem, and knowledge management tool set; 3) training facilities and equipment; and 4) a Web portal for HPC-related information, real-time data, and documentation.

In addition, this project provides the following user support capabilities in conjunction with other projects: 1) a tiered user support structure (HPC service desk) that responds to SNL and tri-lab user requests received via phone, email, Web-based requests, and in-person visits; 2) the Synchronized Account Request Automated Process (SARAPE) tri-lab account provisioning Web-based tool; 3) Web-based, classroom, and one-on-one training; and 4) direct support in utilizing ASC resources.

This project also funds the SNL user support team's involvement in collaborative efforts such as the PSAAP and ACES.

Required Capabilities

R1: HPC service desk resources and infrastructure

R2: Experienced technical staff and support personnel versed in various aspects of HPC, including hardware and software environments, effective application facilitation, and use of all ASC HPC resources

R3: User access mechanisms for HPC resources

R4: Training facilities, documentation, and other information resources

R5: Coordinated tri-lab user support functions

Five-Year Plan

FY15

- Provide user support for SNL and tri-lab ASC computing systems
- Develop expertise in support of next-generation architectures and software environments
- Ramp up to deliver user support for Trinity

FY16

- Provide user support for SNL and tri-lab ASC computing systems
- Deliver user support for next-generation computing environments

FY17

- Provide user support for SNL and tri-lab ASC computing systems
- Continue enhancement of user support resources and capabilities, leveraging trends in new technologies

FY18

- Provide user support for SNL and tri-lab ASC computing systems
- Continue enhancement of user support resources and capabilities, leveraging trends in new technologies

FY19

- Provide user support for SNL and tri-lab ASC computing systems
- Continue enhancement of user support resources and capabilities, leveraging trends in new technologies

Facilities, Networking, and Power (SNL)

The Facilities, Networking, and Power project funds the power and space charges assigned to HPC systems (capacity and file system servers) and long-term hierarchical storage servers (running the HPSS software product). It provides for facilities and personnel to manage installation and removal of computing platforms, file systems, visualization systems, networking equipment, power distribution systems, and cooling systems in support of all computing resources. It also funds major operations contracts such as the ASC Distance Computing (DisCom) WAN.

Facilities professionals have reduced overall operating expenses by minimizing cooling and electrical distribution expenses over the last several years through a comprehensive program of introducing more efficient computer room air conditioning units, using higher voltage electrical source power distribution units, exploring alternative energy sources and conservation mechanisms, which include reducing the volume of chilled water required for cooling and improving air flow in the facility by minimizing obstructions underneath the computer floor. These efforts have been recognized with several SNL-specific and national awards, including three 2011 EStar Awards from the DOE Office of Sustainability Support.

Required Capabilities

R1: Facilities, power, and cooling expertise needed to continue energy conservation practices and inform new facility designs for high-efficiency energy use

R2: Network expertise and continued operational funding for DisCom WAN linking LANL, LLNL, and SNL campuses

Five-Year Plan

FY15

- Analyze networking technologies for next-generation upgrade of DisCom WAN
- Design expansion of Building 725 computer floor space
- Transition 100GE connection to LANL into production state
- Publish definitive evaluation of power requirements comparing identical TLCC2 platforms (one air cooled, one liquid cooled)

FY16

- Configure power and cooling, and manage installation efforts for capability computing platform in expanded NSCC facility
- Provide ongoing operations and maintenance of electrical and mechanical systems for ASC computing facilities

FY17

- Introduce new technologies to DisCom WAN

- Prepare facility for next-generation capacity systems
- Provide ongoing operations and maintenance of electrical and mechanical systems for ASC computing facilities

FY18

- Continue energy saving practices integrating new cooling technologies throughout all facilities
- Provide ongoing operations and maintenance of electrical and mechanical systems for ASC computing facilities

FY19

- Evaluate and document requirements for modern data processing facility to replace aging and inflexible 880 Computing Annex
- Provide ongoing operations and maintenance of electrical and mechanical systems for ASC computing facilities
- Prepare facilities planning for CTS-2

Common Computing Environment (WBS 1.5.5.6)

The goal of the CCE product is to enable a common environment across the tri-labs that was initially deployed on the TLCC systems. The scope of this product includes funded R&D projects to address gap areas identified by the tri-lab technical working groups.

The CCE working groups and projects focus on a common software stack, including but not be limited to, OS software; application development tools; resource management; HPC monitoring and metrics; and common tri-lab environment issues such as configuration management, licenses, WAN access, and multi-realm security.

System Software Deployment for Commodity Technology Systems

TOSS is the software stack that runs on Linux capacity clusters, initiating with TLCC platforms delivered in FY08. The goal of the TOSS project is to increase efficiencies in the ASC tri-lab community with respect to both the utility and the cost of the CCE.

This project delivers a fully functional cluster OS capable of running MPI jobs at scale on CTS hardware. The system must meet CCE requirements for providing a common software environment on CTS hardware across the tri-lab complex, now and into the future.

TOSS provides a complete product with full lifecycle support. Well-defined processes for release management, packaging, quality assurance testing, configuration management, and bug tracking are used to ensure a production-quality software environment can be deployed across the tri-lab in a consistent and manageable fashion.

Required Capabilities

R1: Fully functional cluster OS (kernel, Linux distribution, IB stack and related libraries, and resource manager)

R2: Capable of running MPI jobs at scale on Linux capacity clusters

R3: Full lifecycle support including release management, packaging, QA testing, configuration management, and bug tracking

Five-Year Plan

FY15

- Provide ongoing TOSS software development and support
- Develop/deploy TOSS 2.X (based on RHEL 6.X)
- Develop/deploy TOSS 3.X (based on RHEL 7.X)
- Prepare for deployment of the next generation of ASC CTS systems (CTS-1), which may include software integration and testing for the tri-lab environment
- Continue SLURM support efforts through tri-lab collaboration
- Develop identified collaborative system software tasks, including investigation of new architectures (for example, general-purpose graphics processing units (GPGPUs) and ARM), integration of virtualization, logging/monitoring improvements, and testing infrastructure improvements

FY16

- Develop/deploy TOSS 3.X (based on RHEL 7.X)
- Provide ongoing TOSS software development and support
- Continued assessment and integration of new architecture into TOSS capability

FY17

- Provide ongoing TOSS software development and support
- Continued assessment and integration of new architecture into TOSS capability

FY18

- Provide ongoing TOSS software development and support
- Continued assessment and integration of new architecture into TOSS capability

FY19

- Provide ongoing TOSS software development and support
- Continue assessment and integration of new architecture into TOSS capability

Programming Environment Development/Support for Tri-Lab Systems

The goals of the Programming Environment Project are to enhance productivity of the tri-lab application development teams, operation teams, and analysts. This project achieves these goals by developing and deploying user tools and programming environments to support a variety of applications running on tri-lab HPC resources. It leverages skills across the tri-labs to meet the challenges posed by rapid changes in processor and systems technology and evolving programming models. This project entails software integration, feature enhancements, installation, training, support for vendor provided tools, open source software tools, and lab-developed tools.

Open|SpeedShop (O|SS) is a tri-lab supported open source project that provides a wide range of performance experiments within a single environment. This includes support for PC sampling, inclusive and exclusive execution times for routines, hardware counters, as well as I/O and MPI tracing. O|SS is designed to work on binaries of application without the need to recompile, enabling a clean and easy integration into the development workflow. Transition to the Component-Based Tool Framework (CBTF), which is designed to facilitate easy addition of new capabilities, is currently underway.

Debuggers are another key effort within this project. All three labs are utilizing TotalView as the core debugger. LLNL has the more strategic relationship with Rogue Wave Software and is also doing additional research into other debugging tools (for example, Stack Trace Analysis Tool (STAT)). The approach to subset debugging is key as ASC moves toward larger scales and collects information that helps focus in on the trouble area. Input from all three labs and target applications are required. The working group would like to build a tri-lab debugger capability around the LLNL capability.

The Open Source Contract Maintenance effort provides funding to outside developers who maintain tools and tool infrastructures that are critical for code teams or serve as the basis for internal tools. This funding will be provided to those developers through support contracts administered mainly by LLNL, but each contract includes support for all three laboratories, and all three laboratories in close collaboration provide the technical guidance for the three contracts. This currently includes the tools, Open|SpeedShop, TAU, HPCToolkit, MUST/Vampir, and Valgrind.

The MPI integration/scaling effort is targeted to develop a set of capabilities focusing on supporting scale increases, assessing performance of both MPI and user applications, and providing optimal parameters to users for better MPI performance. The complex multi-socket, multi-core NUMA node architecture of systems such as the CTS-1 mandates such an investigation. A close working relationship with the Open MPI community and other MPI developers is seen as a strategic need.

Required Capabilities

R1: Integration with system software environment to support increased scale

R2: Collaboration with Open Source tool providers and vendors to develop increased integration and leverage software to support increased scale

R3: Ability to support transition to new programming models and system architectures

R4: Support compilation, debugging, performance analysis, and monitoring

Five-Year Plan

FY15

- Validate deliverables on tri-lab contracts for TotalView BIGCAT effort and open-source tools contracts
- Provide user interface enhancements to CBTF
- Provide enhancements and bug fixes to Open MPI/MVAPICH based on tri-lab need; assess MPI performance across many architectures; assess the impact of process and memory binding policies on application performance; and provide results to end users
- Continue development and support efforts for debuggers, performance analysis tools, and MPI as programming models and architectures evolve

FY16

- Continue to integrate programming model and architecture evolution into supported tools

FY17

- Continue to integrate programming model and architecture evolution into supported tools

FY18

- Continue to integrate programming model and architecture evolution into supported tools

FY19

- Continue to integrate programming model and architecture evolution into supported tools

High Performance Computing Environment Integration for Tri-Lab Systems

The HPC Environment Integration project targets the ability to work across sites with minimal transition and access restrictions. Differences in tri-lab security implementation and network restrictions as well resource access and authorization processes have been a hurdle. Efforts target network access infrastructure, cross-realm authentication and resource management and environment standardization. Current efforts include:

- Inter-Site HPC targeted to establish a collaborative environment across sites:
  - Security policies and implementations that allow tri-lab access using identified protocols and technologies and to facilitate easier resource utilization; utilization, where feasible, of home site security apparatus, including authentication equipment (for example, single sign on)
  - User access to tri-lab resources within a specified period of time upon request and based on identified need
  - Support of multiple classes of users, such as architecture and modeling, application, customer service, system administrators with root access, and end users
  - Necessary network bandwidth and latency to cover current and projected requirements
- Establishing a cross-site authentication and resource approval through enhancements to the SARAPE system. This is a Web-based application that allows users within restricted domains to request selected CCE resources to which they are permitted access. It addresses the APIs required to help interface SARAPE with other tools required to manage accounts among the tri-labs. As part of the IHPC deployment, a service catalog will be deployed through which collaborators can view and request accounts and services available in the shared environment.
- The Shared Workspace effort is the infrastructure for promoting collaboration across the laboratories. It currently includes the Gforge server that is housed and managed at SNL.

Required Capabilities

R1: A tri-lab HPC environment that allows ease of cross-site usage; network and security policies; computing resources

R2: Systems and tools in place that collaborative project development, communication, and leveraging of resources

Five-Year Plan

FY15

- Develop increasingly automated functions in SARAPE, including mySARAPE page to allow authenticated users fuller access to account requests and status at all sites

- Migrate CCE tri-lab activities tracking and documentation from gforge to Confluence Wiki or other appropriate repositories
- Continue WC Tool efforts to meet new and/or expanded ASC HQ reporting requirements; address issues in evolving tri-lab computing environments; investigate and develop streamlined ASC HQ reporting tools

FY16

- Deploy changes to site environments based on developed working model
- Continue integration of collaborative tools as required, including SARAPE
- Continue WC Tool efforts to meet new and/or expanded ASC HQ reporting requirements; address issues in evolving tri-lab computing environments

FY17

- Continue changes to site environments based on developed working model
- Continue integration of collaborative tools as required, including SARAPE
- Continue WC Tool efforts to meet new and/or expanded ASC HQ reporting requirements; address issues in evolving tri-lab computing environments

FY18

- Continue changes to site environments based on developed working model
- Continue integration of collaborative tools as required, including SARAPE
- Continue WC Tool efforts to meet new and/or expanded ASC HQ reporting requirements; address issues in evolving tri-lab computing environments

FY19

- Continue changes to site environments based on developed working model
- Continue integration of collaborative tools as required, including SARAPE
- Continue WC Tool efforts to meet new and/or expanded ASC HQ reporting requirements; address issues in evolving tri-lab computing environments

Monitoring and Metrics Integration for Tri-Lab Systems

The Monitoring and Metrics Integration project targets efficient and productive use of HPC systems as well as informed future planning through: 1) effective monitoring of all measurable or reportable conditions on compute platforms, both current and future, that can impact the performance of both applications and throughput on those platforms; and 2) appropriate transformation of monitored information into metrics and transport of those metrics to facilitate their use by system utilities, applications, resource managers, users, system administrators, and management. Integration of information from disparate data sources will enable greater system understanding and response to system conditions.

The project is targeting issues in current systems but with a focus on the increases in both scale and complexity that are expected over the next 5–10 years. With respect to software development, existing tools will be leveraged to the degree possible (for example, Splunk for appropriate data analysis tasks and Web protocols for information transport), and software will be written where gaps exist that cannot be filled by existing software (for example, distributed data services incorporates project-written software at the endpoints and existing transport frameworks).

The project currently focuses on three areas: 1) deployment of developed monitoring infrastructure (for example, Splunk) at all three sites, 2) development of a common scheme for organization, access, and aggregation of monitored information to improve troubleshooting and productivity, and 3) run-time use of platform resource utilization and state information by applications for increased performance.

Required Capabilities

R1: Efficient HPC troubleshooting through platform independent information aggregation and analysis tools

R2: Optimization of resource utilization through effective use of monitored information

R3: Portability of tools and infrastructure both cross-platform and cross-lab

Five-Year Plan

FY15

- Continue development of Splunk tools as needed across the sites (LANL, LLNL, SNL)
- Explore environmental, resource utilization, and health-testing-related HPC monitoring and analysis tools and infrastructure currently being developed/utilized across the tri-lab (for example, Cerebro, Gazebo, LDMS, Lorenz)
- Explore opportunities for collaborative development, integration, and deployment of monitoring and analysis tools and infrastructure across tri-lab ASC resources

FY16

- Continue to evolve monitoring, analysis, and feedback tools as systems and applications evolve
- Test tools in new HPC environments (both production and experimental) to ensure readiness as new platforms arrive

FY17

- Continue to evolve monitoring, analysis, and feedback tools as systems and applications evolve
- Test tools in new HPC environments (both production and experimental) to ensure readiness as new platforms arrive

FY18

- Continue to evolve monitoring, analysis, and feedback tools as systems and applications evolve
- Test tools in new HPC environments (both production and experimental) to ensure readiness as new platforms arrive

FY19

- Continue to evolve monitoring, analysis, and feedback tools as systems and applications evolve
- Test tools in new HPC environments (both production and experimental) to ensure readiness as new platforms arrive

Special Purpose Facilities, Systems, Operations, and Support (WBS 1.5.5.7)

This product provides special purpose high performance computing resources to the DOE community and the necessary support and maintenance of these systems and facilities. This includes special security controls and special purpose facilities in addition to the standard high performance computing operations and support activities necessary to support these resources.

Exploitation (LLNL)

TBD

Required Capabilities

TBD

Five-Year Plan

FY15

- Deliver studies of interest and provide briefings to NNSA

FY16

- TBD

FY17

- TBD

FY18

- TBD

FY19

- TBD

High Performance Computing and Code Support (LLNL)

TBD

Required Capabilities

TBD

Five-Year Plan

FY15

- Provide additional hardware and system administration services for special-purpose networks and clusters
- Adapt physics to codes to support foreign analysis

FY16

- TBD

FY17

- TBD

FY18

- TBD

FY19

- TBD

Special Purpose Computing (SNL)

TBD

Required Capabilities

TBD

Five-Year Plan

FY15

- Continue operations of computing platforms, storage subsystems, and data archive
- Expand capabilities for remote access, data transfer, and remote graphics services
- Implement COI data protection mechanisms
- Initiate procurement activity for NSCC computing resources
- Complete security plan and usage models
- Fully deploy a user support process for NSCC

FY16

- TBD

FY17

- TBD

FY18

- TBD

FY19

- TBD

Appendix A. Glossary

ACES	New Mexico Alliance for Computing at Extreme Scale
ADEPT	Applications Development Environment and Performance Team
AMD	Advanced Micro Devices
AMR	Adaptive Mesh Refinement
AMT	Asynchronous Many Task
ANL	Argonne National Laboratory
API	Application Programming Interface
ASC	Advanced Simulation and Computing
ASCR	Advanced Scientific Computing Research (DOE/SC)
ATCC	Advanced Technology Computing Campaign
ATDM	Advanced Technology Development and Mitigation
ATS	Advanced Technology System
BG/Q	BlueGene/Q
CBTF	Component-Based Tool Framework
CCC	Capability Computing Campaign
CCE	Common Computing Environment
CD	Critical Decision
CORAL	Collaboration of Oak Ridge, Argonne, and Livermore
CPU	Central Processing Unit
CSSE	Computational Systems and Software Environment
CTBF	Component-Based Tool Framework
CTS	Commodity Technology System
CZ	Collaboration Zone
D&E	Development and Engineering
DisCom	Distance Computing
DOE	Department of Energy
DRAM	Dynamic Random Access Memory

DSW	Directed Stockpile Work
EDR	Enhanced Data Rate
FOUS	Facility Operations and User Support
GPFS	General Parallel File System
GPU	Graphics Processing Unit
HMC	Hybrid Memory Cube
HPC	High Performance Computing
HPSS	High-Performance Storage System
HQ	ASC Headquarters
I/O	Input/Output
IB	InfiniBand
IC	Integrated Codes
iHPC	Inter-Site High Performance Computing
ITIL	Information Technology Infrastructure Library
LAN	Local Area Network
LANL	Los Alamos National Laboratory
LC	Livermore Computing
LDMS	Lightweight Distributed Metric Service
LFLR	Local Failure Local Recovery
LLNL	Lawrence Livermore National Laboratory
LWK	Lightweight Kernel
MIC	Many Integrated Core
Mini-App	Mini Application
MOU	Memorandum of Understanding
MPI	Message Passing Interface
NAS	Network-Attached Storage
NFRSC	National Energy Research Scientific Computing Center
NFS	Network File System
NGBB	Next-Generation Backbone
NNSA	National Nuclear Security Administration
NRE	Non-Recurring Engineering

NSCC	National Security Computing Center
NUMA	Non-Uniform Memory Access
NVRAM	Non-Volatile Random Access Memory
nWBS	National Work Breakdown Structure
NWC	Nuclear Weapons Complex
OS	Operating System
O|SS	Open|SpeedShop
ORNL	Oak Ridge National Laboratory
OTP	One-Time Password
PDE	Partial Differential Equation
PEM	Physics and Engineering Models
PI	Principal Investigator
PLFS	Parallel Log File System
PSAAP	Predictive Science Academic Alliance Program
PSP	Predictive Science Panel
QA	Quality Assurance
QASPR	Qualification Alternatives to the Sandia Pulsed Reactor
R&D	Research and Development
RFP	Request for Proposal
RHEL	Red Hat Enterprise Linux
RZ	Restricted Zone
SAIC	Science Applications International Corporation
SAN	Storage Area Network
SARAPE	Synchronized Account Request Automated Process
SC	Department of Energy's Office of Science
SCC	Simulation and Computing Complex
SCF	Secure Computing Facility
SIO	Scalable Input/Output
SLURM	Simple Linux Utility for Resource Management
SNL	Sandia National Laboratories
SQL	Structured Query Language

SRAM	Static Random Access Memory
SSP	Stockpile Stewardship Program
SST	Structural Simulation Toolkit
STAT	Stack Trace Analysis Tool
TLCC	Tri-Lab Linux Capacity Cluster
TOSS	Tripod Operating System Software
UQ	Uncertainty Quantification
V&V	Verification and Validation
VM	Virtual Machine
VTK	Visualization Toolkit
WAN	Wide Area Network
WBS	Work Breakdown Structure
WC	Workload Characterization